



**A NONPARAMETRIC APPROACH TO
SEGMENTATION OF LADAR IMAGES**

DISSERTATION

Eric A. Buschelman, Captain, USAF
AFIT/DEE/ENG/12-07

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT/DEE/ENG/12-07

A NONPARAMETRIC APPROACH TO SEGMENTATION OF LADAR IMAGES

DISSERTATION

Presented to the Faculty

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

Eric A. Buschelman, B.S.E.E., M.S.E.E.

Captain, USAF


December 2012

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

A NONPARAMETRIC APPROACH TO SEGMENTATION OF LADAR IMAGES

Eric A. Buschelman, B.S.E.E., M.S.E.E.
Captain, USAF


Approved:


Richard K. Martin, PhD (Chairman)

31 Aug 2012
Date



Stephen C. Cain, PhD (Member)

21 Sept 2012
Date


Christine Schubert Kabban, PhD (Member)

5 Sep 2012
Date

Accepted:


M. U. Thomas
Dean, Graduate School of Engineering and Management

10 Oct 2012
Date

Abstract

This research project explores full-waveform analysis of three-dimensional “flash” laser radar (ladar) signals for the purpose of segmenting and classifying ladar images. A novel nonparametric range data segmentation designed for experimentally collected ladar data is presented, and yields results superior to conventional edge and boundary detection methods. The completeness of the segmentation enables calculation of additional target characteristics, such as the orientation angle and reflectance of the illuminated objects. Although a cross-correlation ranging algorithm is used to process the collected ladar data, the additional information gained from the segmentation routine is a new contribution to understanding the range of capabilities of the ladar camera used throughout this research.

In addition to the nonparametric segmentation, a widely accepted ladar waveform model is updated to include the effects of optical backscattering from elongated targets. The effective gain of the waveform model accounts for target incidence angle, and shows agreement with independently published research. A successful waveform analysis is achieved by estimating signal amplitude, range, and bias, and the results are compared to other accepted methods.

Additionally, local statistics of the waveform and object parameters are used in an unsupervised nonparametric classification routine. Local statistics have been used frequently in image analysis to differentiate between both objects and regions, and this approach is attempted here. Feature selection of both the segment point data and higher-order statistics emphasizes intrinsic material properties, given the emphasis on classifying similar objects together. Two-sample rank tests of location and dispersion form the core of the classification method.

AFIT/DEE/ENG/12-07

For Mother and Father

Acknowledgements

This latest AFIT journey really began for me back in 2005, when my undergraduate advisor at the University of Kansas, Dr. Ken Demarest, recommended me for the AFIT Master's program as a direct accession 2d Lt. Appropriately, I've now earned graduate degrees in my two worst undergraduate subjects: electromagnetism and statistics. Without his encouragement, who knows what my Air Force career would look like now.

I am grateful also for the encouragement of those officers who helped push me over the edge into the PhD program back in 2009: Col (r) Rick Buschelman, Col Todd Krueger, Col Kurt Hall, Lt Col Mark Skouson, and Capt Rob Johnson. Their advice was – and remains – invaluable.

These last three years at AFIT have been filled with individuals who have shown me great friendship, and I am indebted to them. Fellow members of my section, who in their wisdom chose me to be the Section Leader for the last year, have been mentors, leaders, and role models for me: ATIS, Smiley, Shock, Zilla, Fabio, 9-Point, Hummer, Tiko, Dog, Spot, Bull, DB and Vap. Huge thanks also go to Maj Brian Neff, who was of great assistance during the data collection phase of my research; I would have been a lost student without his resourcefulness. Capt Jason Vap also helped me remodel my bathroom at the worst possible time imaginable: 2 months before my defense.

Special thanks go to Dr Rusty Baldwin, Chaplain (LtCol) Srode and the Wright-Patt Chaplain's office for providing weekly Mass at AFIT for the better part of the last 6 years. Thanks be to God for this great blessing!

I spent many afternoons running with members of the AFIT Intramural Running Club, and had the good fortune to run with the "Legends" during their final years at AFIT: Col David Cribb, LtCol Doug Wickert, LtCol Amie Stryker, LtCol Jeff McGuirk, Lt Col Brad Christiansen, and Maj Craig Larsen. We all wanted to keep the good thing you started going. In the more recent years, in addition to members of my own section, I was happy to be team mates with CDR Neal "Dirty" Kraft, Lt Alex Hollenbeck, Capt Randy Cicale, Capt Jonathan Campbell, and Capt Cindy Finke, who took over as team captain in the

last year. Without the help of Mr Mike McInturf in the Intramural Sports office, none of it would have happened. Running IM's was often the highlight of my week, and I will miss the competition as I move on.

I was fortunate to have in Dr Rick Martin an advisor who could manage my chronic indecisiveness concerning research direction. In our 3 years together, he was always able to see through the fog in which I had surrounded myself, identify what I was trying to do, and pull me out of it. This was a regular occurrence, and it was a lifesaver every time. My research exists because of his firm guidance. Thank you.

My other committee members were also instrumental in shaping my time at AFIT. Dr Cain's singularly best piece of advice (of many) was to "get a laptop, and put MATLAB on it." He was right. Dr Schubert-Kabban saw me through from my early days as dyed-in-the-wool electrical engineer, all the way to a full-blown wannabe statistician: just look at the title of the dissertation for proof of that. I truly value her contribution to my education and worldview, as well as our many cordial talks throughout the years.

To all the Official Members of the Unofficial AFIT Home Brew Club – Fabio, Bull, Hummer, 9-Point, Zilla, Raj, Jorge – cheers.

At last I thank my family, beginning with my parents. Our Air Force careers couldn't have been any more different, but I am thankful for the legacy, now in its third generation, that I am part of. To my siblings and in-laws, who never really pushed for more details when they asked me: "so what is it you're studying?" but were always supportive. The great joys in my life who are my three wonderful, curious, and precious children: I love you and yes, Daddy is a doctor now. To my adoring and supportive wife of seven years: you are my best friend. Thank you for sharing your life with me.

Omnia possum in eo qui me confortat (Phil 4:13). Totus tuus ego sum, et omnia mea tua sunt.

Eric A. Buschelman

Table of Contents

| | Page |
|--|-------|
| Abstract | iv |
| Acknowledgements | vi |
| List of Figures | xii |
| List of Tables | xvi |
| List of Abbreviations | xvii |
| List of Symbols | xviii |
| I. Introduction | 1 |
| 1.1 Motivation | 2 |
| 1.1.1 Ladar Signal Model | 3 |
| 1.1.2 Ladar Images | 3 |
| 1.2 Research Contributions | 4 |
| 1.2.1 Nonparametric Ladar Data Segmentation (Chapter IV) | 4 |
| 1.2.2 Backscatter Reflectance Waveform Model (Chapter V) | 4 |
| 1.2.3 Ladar Feature Analysis (Chapter VI) | 5 |
| 1.2.4 Nonparametric Classification for Ladar Images (Chapter VII) | 5 |
| 1.3 Organization | 5 |
| II. Background | 7 |
| 2.1 Full Waveform Lasers | 7 |
| 2.2 Standard Pulse Models | 9 |
| 2.3 Image Segmentation | 9 |
| 2.4 Range Data Segmentation | 12 |
| 2.5 Classification Techniques | 13 |
| 2.6 Nonparametric Statistics | 15 |
| 2.7 Lower Bounds on Estimator Performance | 16 |
| III. Experimental Data Collection and Waveform Analysis | 18 |
| 3.1 Ladar Camera Properties | 18 |
| 3.2 Target Data Collection | 19 |
| 3.3 Ladar Data Processing | 20 |
| 3.3.1 Raw Data Processing | 21 |
| 3.3.2 Compound Poisson Model | 22 |
| 3.3.3 Temporal Pulse Model | 24 |
| 3.3.4 Range Cross-Correlation | 26 |

| | Page |
|--|------|
| 3.4 Summary | 29 |
| IV. Nonparametric Ladar Image Segmentation | 30 |
| 4.1 Nonparametric Range Data Segmentation | 32 |
| 4.1.1 Kernel Density Estimation | 33 |
| 4.1.2 Density Range Bins | 34 |
| 4.1.3 Boundary Tracing and Segmentation | 36 |
| 4.1.4 Examples | 39 |
| 4.1.5 Summary | 41 |
| 4.2 Iterative Region Growing Segmentation | 42 |
| 4.2.1 Segmentation Algorithm | 44 |
| 4.2.2 Results | 45 |
| 4.2.3 Summary | 49 |
| 4.3 Third-Phase Ladar Image Analysis | 49 |
| 4.3.1 Plane Fitting | 50 |
| 4.3.2 Material Reflectance | 53 |
| 4.3.3 Examples | 55 |
| 4.3.4 Summary | 57 |
| 4.4 Chapter Summary | 58 |
| V. Backscatter Reflectance Waveform Model | 59 |
| 5.1 Backscatter Reflectance | 59 |
| 5.1.1 Waveform Parameter Estimates | 65 |
| 5.2 Minimum Error Bounds | 69 |
| 5.2.1 Likelihood Based on Product | 70 |
| 5.2.2 Likelihood Based on Sum | 74 |
| 5.2.3 Estimator Performance | 75 |
| 5.3 Other Least-Error Bounds | 78 |
| 5.3.1 Bobrovksy-Zakai Bound | 79 |
| 5.3.2 Barankin Bound | 81 |
| 5.3.3 Bayesian Abel Bound | 82 |
| 5.3.4 Simulation | 84 |
| 5.3.5 Discussion | 85 |
| 5.4 BRWM Verification | 85 |
| 5.4.1 Incidence Angle Estimation | 85 |
| 5.4.2 Amplitude Variation | 86 |
| 5.5 Chapter Summary | 88 |

| | Page |
|---|------|
| VI. Ladar Feature Analysis | 89 |
| 6.1 Feature Extraction | 89 |
| 6.1.1 Image Statistics | 90 |
| 6.1.2 Sample Density Estimation | 91 |
| 6.2 Feature Selection Examples | 91 |
| 6.2.1 Multi-Facet Board Simulation | 92 |
| 6.2.2 Field Experiments | 93 |
| 6.3 Chapter Summary | 96 |
| VII. Nonparametric Classification for Ladar Images | 102 |
| 7.1 Classification Tests and Features | 102 |
| 7.1.1 Wilcoxon Rank Sum Test | 103 |
| 7.1.2 Ansari-Bradley Test | 104 |
| 7.1.3 Feature Selection | 105 |
| 7.2 Classification Methodology | 106 |
| 7.3 Examples | 107 |
| 7.4 Summary | 110 |
| VIII. Conclusions | 112 |
| 8.1 Chapter Summaries | 112 |
| 8.2 Summary of Contributions | 113 |
| 8.2.1 KDE-Based Range Data Segmentation | 113 |
| 8.2.2 Third-Phase Ladar Image Analysis | 113 |
| 8.2.3 Backscatter Reflectance Waveform Model | 114 |
| 8.2.4 Image Classification using Nonparametric Tests | 114 |
| 8.3 Future Research Possibilities | 114 |
| 8.3.1 Integration of De-blurring and Range Data Segmentation | 115 |
| 8.3.2 Anomaly Detection using 3D Flash Ladar | 115 |
| 8.3.3 Covariance of 3D Flash Ladar Systems | 115 |
| 8.3.4 Full Waveform Analysis of Multiple Returns | 115 |
| A. Derivation of Cramer-Rao Bound | 117 |
| 1.1 Preamble | 117 |
| 1.2 CRB | 118 |
| B. Derivation of Bayesian Abel Bound | 125 |
| 2.1 Deterministic CRB | 125 |
| 2.2 Prior Information | 127 |
| 2.3 Bobrovsky-Zakai Bound | 127 |

| | Page |
|-------------------------------|------|
| 2.4 Bayesian Abel Bound | 129 |
| Bibliography | 131 |
| Vita | 139 |

List of Figures

| Figure | | Page |
|--------|--|------|
| 3.1. | (a) ASC Portable Ladar Camera mounted to machined brass plate, in storage box. (b) Camera deployed for collecting images | 19 |
| 3.2. | Representative photo of data collection target. | 20 |
| 3.3. | Birds-eye view cartoon of data collection targets. Board 2 is titled with respect to the camera throughout the experiment. | 21 |
| 3.4. | Sample Gaussian pulse model from Eq. (3.8). | 25 |
| 3.5. | Matched filtering comparison of parametric and non-parametric method. | 28 |
| 3.6. | Correlation coefficients and p-values for matched filter output. | 29 |
| 4.1. | Sample “first-phase” principal range image. | 30 |
| 4.2. | Sample “second-phase” 2D intensity image. | 31 |
| 4.3. | Sample range density histograms | 33 |
| 4.4. | Sample range density using Gaussian kernel. Bandwidth values: (a) 0.17 m (b) 0.35 m (c) 0.70 m | 35 |
| 4.5. | Sample range density estimate, calculated using Gaussian kernel. Bandwidth: 0.17 m | 36 |
| 4.6. | Sample range surface image. | 38 |
| 4.7. | $n = 1$ range density bin images. | 38 |
| 4.8. | Composite boundaries \mathfrak{B} , overlaid on 2nd phase intensity image. Density bandwidth: 0.17 m | 38 |
| 4.9. | Composite boundaries \mathfrak{B} , overlaid on 2nd phase intensity image. Density bandwidth: 0.35 m | 39 |
| 4.10. | Principal range image of 6 Feb 12 range data, trials 1–4. | 40 |
| 4.11. | Density (BW = 0.35 m) and segmentation of 6 Feb 12 data, target 1. | 40 |
| 4.12. | Density (BW = 0.35 m) and segmentation of 6 Feb 12 data, target 2. | 41 |

| Figure | Page |
|---|------|
| 4.13. Density (BW = 0.35 m) and segmentation of 6 Feb 12 data, target 3..... | 42 |
| 4.14. Density (BW = 0.35 m) and segmentation of 6 Feb 12 data, target 4..... | 43 |
| 4.15. Canny edge detection of 6 Feb 12 range data, trials 1–4. | 43 |
| 4.16. Segmentation using iterative SRG on 6 Feb 12 range data, trials 1–4. | 46 |
| 4.17. Segmentation using nonparametric range KDE, on 6 Feb 12 range data, trials 1–4..... | 47 |
| 4.18. System diagram, showing target aspect angle with respect to laser boresight..... | 52 |
| 4.19. Sample 3rd phase image, showing fit-plane inclination angle. Angles are measured in degrees from the z -axis. | 52 |
| 4.20. Edges of Fig. 4.19 using Canny edge filter..... | 53 |
| 4.21. System gain estimate $\hat{\mathfrak{G}}$ density graphs. (a) Untrimmed histogram (b) 10% trimmed mean histogram (c) Untrimmed density and (d) 10% trimmed mean density. | 55 |
| 4.22. Relative α estimate. | 55 |
| 4.23. Third-phase Images of Trial 1 data | 56 |
| 4.24. Third-phase Images of Trial 2 data..... | 57 |
| 4.25. Third-phase Images of Trial 3 data..... | 57 |
| 4.26. Third-phase Images of Trial 4 data..... | 57 |
| 5.1. Contours of effective gain $G_{eff} \propto (\mathfrak{G}\alpha \cos \theta_i) / (\sqrt{2\pi}\sigma_R)$. Reflectance parameter $\alpha = 0.8$ | 61 |
| 5.2. Sample Gaussian pulse models from Eq. (3.8) (Standard) and Eq. (5.5) (BRW). Inclination angle $\theta_i = 25^\circ$, and $\alpha = 0.7$ | 62 |
| 5.3. Normalized peak pulse amplitude of standard and BRW pulse models. Pulse spreading effects are ignored. | 63 |

| Figure | Page |
|--|------|
| 5.4. Signal pulse energy with respect to incidence angle. Comparison is made between BRW and Standard models of pulse amplitude..... | 64 |
| 5.5. Time-amplitude plots of the elements of \mathbf{J} . (a) J_{11} (b) J_{22} (c) J_{33} (d) J_{13} (e) J_{12} (f) J_{23} . Parameters for the plots are shown in Tab. 5.1. | 73 |
| 5.6. CRB of range matched-filter estimate with (a) $B = 100$, and (b) $B = 1000$ | 76 |
| 5.7. CRB of gain estimate with (a) $B = 100$, and (b) $B = 1000$ | 78 |
| 5.8. CRB of bias estimate \hat{B} with (a) $B = 100$, and (b) $B = 1000$ | 79 |
| 5.9. Convergence of Bayesian bounds | 84 |
| 5.10. Ladar system diagram | 86 |
| 5.11. Measured mean peak amplitude of target boards with respect to laser incidence angle. | 87 |
| 6.1. Simulation target (a) range profile, (b) intensity profile..... | 94 |
| 6.2. Peak amplitude of target boards with respect to aspect angle. Truth reflectance $\alpha = 0.7$ | 94 |
| 6.3. Color photos and intensity images of the two target setups. The plywood board in setup 2 is tilted at $\sim 15^\circ$ with respect to the vertical board. | 97 |
| 6.4. Average waveforms reflected from the white and black boards, taken from Fig 6.3(c). The ratio of the peak intensity values is 0.62. | 98 |
| 6.5. Bootstrapped mean value for parameters of hand-labeled objects from both target setups (1000 samples). Intensity is measured in photon counts, and pulse width is measured in seconds. | 99 |
| 6.6. Bootstrapped mean range-to-targets. | 100 |
| 6.7. Bootstrapped means of local statistics (1000 samples). Standard deviation computed using 5×5 window, and skewness computed using 7×7 window. | 101 |

| | | |
|------|---|-----|
| 7.1. | Histograms of local statistic $sk\left(\hat{\mathfrak{S}}_{\alpha}\right)$ for each segmented region. Window Size: 8 pixels. Dataset: 6 Feb '12, Trial 4..... | 109 |
| 7.2. | Patch segmentation and system gain map of 6 Feb 12 data, target 4..... | 109 |

List of Tables

| Table | Page |
|--|------|
| 3.1. 3D Ladar Camera Parameters. | 19 |
| 3.2. Per-pixel Bias Variance (counts) | 22 |
| 3.3. Waveform parameters for the sample plots of Fig. 3.6 and Fig. 3.6. | 29 |
| 4.1. Area & Classification Rates (CRs) of Fixed Board Segmentation. | 47 |
| 4.2. Area & Classification Rates (CRs) of Angled Board Segmentation. | 48 |
| 4.3. Area & Classification Rates (CRs) of Fence Segmentation. | 48 |
| 4.4. Calculated Inclination Angles. | 56 |
| 5.1. Waveform parameters for the sample plots of Fig. 5.5. | 72 |
| 5.2. Waveform parameters for the sample plots of Fig. 5.6. | 76 |
| 6.1. Simulation Waveform parameters. | 93 |
| 6.2. Simulation. Inclination angle statistics. Angles are in degrees. | 93 |
| 6.3. Simulation. Reflectance parameter statistics. | 94 |
| 7.1. Classification Features | 106 |
| 7.2. Wilcoxon Test Outcomes. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8 | 108 |
| 7.3. Ansari-Bradley Test Outcomes. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8 | 108 |
| 7.4. Classification Results. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8. | 108 |
| 7.5. Wilcoxon Test Outcomes. $ku(I)$, Window Size: 15 | 110 |
| 7.6. Ansari-Bradley Test Outcomes. $ku(I)$, Window Size: 15 | 110 |
| 7.7. Classification Results. $ku(I)$, Window Size: 15 | 110 |

List of Abbreviations

| Abbreviation | Page |
|--------------|---|
| ALS | Airborne laser scanning 8 |
| ROC | Receiver Operating Characteristic 8 |
| RGB | Red-Green-Blue 12 |
| PCA | Principal Component Analysis 12 |
| AUC | Area Under the ROC Curve 14 |
| KDE | Kernel Density Estimation 15 |
| CRB | Cramér-Rao Bound 16 |
| BRB | Barankin Bound 16 |
| BCRB | Bayesian Cramér-Rao Bound 16 |
| BZB | Bobrovksy-Zakai Bound 17 |
| RMB | Reuven-Messer Bound 17 |
| BAB | Bayesian Abel Bound 17 |
| ASC | Advanced Scientific Concepts 18 |
| FOV | Field-of-view 18 |
| LOS | Line-of-sight 18 |
| FWHM | Full-Width Half-Maximum 25 |
| AWGN | Additive White Gaussian Noise 26 |
| SRG | Seeded region growing 42 |
| BRDF | Bi-directional Reflectance Distribution Function 60 |
| BRW | Backscatter Reflectance Waveform 62 |
| MLE | Maximum Likelihood Estimation 65 |
| WRST | Wilcoxon Rank Sum Test 103 |
| CDF | Continuous Distribution Function 103 |

List of Symbols

| Symbol | | Page |
|-----------------|---|------|
| F_0 | Focal length | 18 |
| τ_R | Range resolution | 20 |
| σ_w | Pulse length | 20 |
| θ_i | Incidence angle | 20 |
| R | Principal target range | 25 |
| G_{eff} | Effective gain | 25 |
| B | Noise bias | 25 |
| τ_G | Gaussian pulse FWHM | 25 |
| h_n | KDE smoothing function | 34 |
| \mathcal{R}_n | Set of regions in range bin n | 37 |
| r_k | k th contiguous region of \mathcal{R}_n | 37 |
| R_Σ | Total number of segmented regions | 37 |
| \mathcal{B}_n | Set of boundaries of \mathcal{R}_n | 37 |
| b_k | Boundary of r_k | 37 |
| \mathfrak{B} | Composite boundary set of all segments | 37 |
| R_{med} | Median range of a segment | 50 |
| d_{trans} | Transverse scene width | 50 |
| \mathcal{D}_k | Data points in r_k | 51 |
| \vec{n} | Surface normal vector | 51 |
| \vec{e}_z | z -axis unit vector | 51 |
| \mathfrak{G} | Ladar system gain | 54 |
| α | Wavelength-dependent surface reflectance | 54 |

| Symbol | | Page |
|------------------------|---|------|
| $\widehat{\theta_i^k}$ | Incidence angle θ_i of k th region | 54 |
| L | Radiance | 60 |
| J | Fisher Information Matrix | 70 |
| A | Vector of waveform parameters | 70 |

I. Introduction

The emerging technological development of three-dimensional (angle-angle-time) laser radar is made possible by a focal plane array capable of collecting a series of two-dimensional images of a remote scene from one laser pulse, vastly increasing the amount of available waveform data [65] [94]. The so-called “flash” ladar collects an entire frame of ladar data in just a single pulse, creating the potential for significant increases in the amount of usable sensor data.

Data collection at higher rates encourages use of this new technology for collision avoidance, object tracking, and target recognition, in addition to applications within cartography, forestry, astronomy and military domains [102]. At the same time, the demand for information gleaned from these remote sensing devices has grown exponentially, prompting researchers to develop techniques for greater automation of data analysis in order to reduce the human workload. The newly developed processes are often rooted in familiar image processing techniques, which are used to parse the data and ultimately to select and discriminate objects of interest from those deemed to be in the “background.”

Recent technological advances permit the collection and storage of the full signal waveform detected by the ladar sensor. Earlier generations of scanning ladar systems stored only the “point cloud” data - range (time to target) and signal amplitude (peak intensity). With the full-waveform available, additional range depth information at each pixel of the scene may be extracted using signal processing. Adjustments made to common ranging algorithms include detection of the peak pulse, “first pulse” or “last pulse” in order to properly characterize the range estimate of the target object. Precise ranging has been the subject of many studies in recent years, and bounds quantifying the range estimates of laser radar systems have recently been proposed [45] [65].

In addition, the precision measurement capabilities of laser radar, especially of scanning systems, are well-suited for imaging applications, and have been used in both airborne and ground settings to great effect. Within the context of defense-oriented applications, classification of objects within the image, whether for mapping, navigation, targeting, identification, or other purposes, can be quite literally a matter of life and death. Assessing the confidence, or likewise uncertainty, in the classification decisions is thus an objective of no little importance.

Flash ladar is a developing technology that builds a *3D* scene by capturing successive *2D* photo-intensity images. The high data rate of the circuitry enables a multitude of data-intensive applications, such as topological mapping, real-time *3D* object tracking, and rapid battlespace visualization. The compact size of the flash ladar apparatus make it an appealing choice for space- or weight-limited functions, including for use on mobile autonomous-navigation robots [101] and on airborne platforms [24]. Motivation for this research is established by discussing the challenges of exploiting full-waveform ladar data for use in image applications. An overview of the specific research contributions is also provided, and is followed by presentation of the document organization.

1.1 Motivation

The original research motivation grew from an interest in the measurement of pulse-expansion in the detected ladar waveform. It has been suggested that additional information about the illuminated target could be extracted from the pulse width measurement, which would represent a potentially significant gain in understanding which was made possible by the flash ladar technology. In addition to an improved range to target calculation, this finding would lead to better modeling scenarios and image analysis. The present research effort grew out of this goal. As the research progressed, emphasis shifted from strict pulse shape analysis to a more general treatment of the features and objects within the ladar image. Exploiting the complete ladar cuboid made it possible to obtain new information about the target without relying heavily on assumptions of pulse shape. The need to segment the image arose out of a consideration for computer and machine vision applications,

and is further supplemented by a nonparametric classification analysis.

1.1.1 Ladar Signal Model.

As one of the primary applications of flash ladar technology is ranging, it is appropriate to consider methods to improve the signal model. This is first done by recognizing the detection of photoelectrons by the ladar detector array as a compound Poisson process with a time-varying mean. The parameter of the Poisson process is the shape of the laser pulse, which has been modeled throughout as a Gaussian pulse. Accounting for the backscatter reflectance of the illuminated targets is the next improvement made to the waveform model. The contributions of Johnson [45] in the area of pulse shape were directed at unresolved point targets occupying a single pixel; since the present research is occupied with resolved, non-point targets, Johnson’s analysis does not immediately transfer. For that reason it was necessary to update the Gaussian waveform model to account for reflection from a resolved, angled target.

1.1.2 Ladar Images.

The complete dataset generated by the 3D flash ladar camera can be deconstructed in two distinct ways. Breaking down the dataset to the pixel level is the full-waveform approach, which measures the photo-electron count at each volume element (voxel) in the angle-angle-time cuboid. Treating the data in this manner is similar to the operation of scanning full-waveform ladar systems, which builds a three-dimensional cuboid by rastering a laser pencil beam across the scene, dwelling at each point.

The alternative breakdown is to view the two-dimensional photo-count intensity image at each time sample in the angle-angle-time cuboid. This visualization mimics the actual operation of the camera, which illuminates (“flashes”) the entire scene with a single broad, short laser pulse. Each pixel in the 2D time slice contains the photo-electron count corresponding to the object located at that range and spatial location. The ratio between photo-electron counts and “camera-counts” of the camera detector array is roughly one-to-one, although this has not been rigorously verified. Various atmospheric blurring

phenomenon, whose effects are measureable in the detected waveform, have been described at length in other publications, such as [66], [68] and [69], and are not discussed further here.

1.2 Research Contributions

1.2.1 Nonparametric Ladar Data Segmentation (Chapter IV).

In this chapter, the proposed ladar segmentation method is introduced and demonstrated. A short discussion of the nature of laser radar images is presented, and this motivates the necessity of organizing ladar images into distinct “phases” of imagery. The segmentation method is initialized using nonparametric probability density estimation. The resulting probability density is sliced piecewise into probability range bins, and the dominant object regions in each slice are traced and labeled. Plane fitting of each region is accomplished using principal components analysis, from which the local surface inclination angle with respect to the ladar boresight is calculated. Finally, the relative material reflectance of the segmented objects is estimated. A discussion of the strengths and weaknesses of the segmentation method is also presented.

1.2.2 Backscatter Reflectance Waveform Model (Chapter V).

This chapter emphasizes the importance of accounting for the backscatter angle when modeling waveforms reflected from extended surfaces. The third-phase imagery of Chapter IV makes it possible to calculate the incidence angle of the illuminated targets. The waveform model of Chapter V is generalized to include the concept of effective gain, which takes into account the reduced amplitude of such a backscattered signal. Implications to other waveform parameter estimates are discussed, including range and pulse shape. The chapter concludes with a lengthy discussion of estimator performance, which is deemed satisfactory when evaluated using the Cramér-Rao Bound.

1.2.3 Ladar Feature Analysis (Chapter VI).

In this chapter the rationale behind the feature selection methodology is examined. Point data and locals statistics of the image data in each segmented region are analyzed for suitability in a classification routine. The material reflectance and measured intensity, and higher-order local statistics thereof, are selected as the most likely candidates for classification. Calculating higher-order statistics often requires a minimum sample size to achieve confidence in the distribution, and so bootstrap resampling is used to overcome any sample size deficiencies. The feature extraction method is tested using both simulated and experimentally collected data, proving the concept feasibility. The experimental outcomes have been published in [12].

1.2.4 Nonparametric Classification for Ladar Images (Chapter VII).

A classification methodology based on nonparametric two-sample tests is described and implemented for use on segmented 3D ladar data. Nonparametric tests have been little used in the ladar imaging community, and this contribution presents classification techniques from other fields (such as medical imaging) for consideration. Collective data points from each of the segmented regions are taken pairwise, and the Wilcoxon rank sum test is used to test for location (median) of the distribution, and the Ansari-Bradley test is used to test for dispersion (variance). Necessary assumptions about the data (such as independence) are taken into consideration, and potential shortcomings of these assumptions, as well as the implementation of the tests themselves, are given a thorough treatment. Results of the classification methodology are mixed, but the technique is promising nonetheless.

1.3 Organization

The dissertation is organized in the following manner. Relevant background material concerning full-waveform ladar, statistical image processing, as well as a through discussion on least-error bounds is presented in Chapter II. Specifics of the data collection, and the basic ranging algorithm used on the data, are reviewed in Chapter III. The non-parametric

approach to image segmentation for this data, which is the first major contribution, is discussed in Chapter [IV](#). The waveform model is updated in Chapter [V](#) to include backscatter effects. Chapter [VI](#) discusses additional feature analysis of the data that can be gleaned from the segmentation. The novel nonparametric classification methodology is put forth in Chapter [VII](#). The document concludes in Chapter [VIII](#) with an overview of each chapter's content, a summary of relevant contributions, and suggestions for future research opportunities.

II. Background

A broad overview of the topics addressed by the research project are presented in this chapter. These topics include: operation and use of full-waveform laser radar systems, especially as used at AFIT and at other institutions; common waveform models and their performance, deficiencies, and suitability; image segmentation and classification techniques, including the appropriate statistical methods; and suitability of the Cramér-Rao and other least-error bounds for evaluating estimator performance.

The chapter is organized as follows: Sec. 2.1 discusses the short history and operation of full-waveform laser radars, addressing both scanning and flash systems. Laser pulse models are reviewed in Sec. 2.2. Different methods of image segmentation are described in Sec. 2.3, while range data segmentation is reviewed in Sec. 2.4. Varying statistical approaches to classification are described in Sec. 2.5. Sec. 2.6 reviews the contemporary uses of nonparametric statistics within image processing. A complete discussion of least-error bounds is given in Sec. 2.7.

2.1 Full Waveform Lasers

For many years, laser radar systems gathered two pieces of information about the sensed target: range and amplitude. The unique properties of operation at optical wavelengths means that lidar has the potential to gather highly precise range measurements. Recent technological advances in all types of lidar systems allow for digitization and recording of the full-waveform, i.e. the received signal of the reflected laser pulse. As a result, a finer level of detail about the scene can be extracted from the received data. Still, this full-waveform capture was often found only in scanning laser radar systems. The recent proliferation of 3D flash lidar cameras allows the capture of entire 3D lidar data frame by a single lidar pulse. The enabling technology, the 3D focal plane array, incorporates rows and columns of pixels, each of which accurately and independently counts the time to the target. Operating in the same manner as a 2D camera, a broad short lidar pulse is “flashed” to illuminate the target. Upon the return of the pulse, it passes through the camera optics and focuses

the portion of the image onto the corresponding pixel in the detector array [93] [94].

Copious amounts of research have been dedicated to exploiting the full-waveform signal in order to obtain improved range estimates. Persson, working mainly with airborne laser scanning (ALS) systems, has demonstrated the effectiveness of extracting additional range information (e.g. multiple returns) from the digitized laser waveform [74]. The points classified as non-ground (i.e. vegetation) were found to have both a wider pulse and lower intensity as compared to the ground-classified points.

Reitberger *et al.* used full-waveform decomposition of ALS data to perform segmentation of individual trees [78], [79]. Gross *et al.* use the features derived from full-waveform decomposition for the purposes of segmenting trees; the performance is evaluated using receiver operating characteristic ROC analysis [34]. The ROC curve shows a higher-than-desireable false-alarm rate for this method. Wagner compares the Average Square Difference Function (ASDF) method to more classical pulse-detection methods of full-waveform analysis, such as echo-detection and time delay estimation, with comparable results [103].

In an investigation of the laser system properties, McMahon has demonstrated improved range estimation of a 3-D flash lidar system using blind deconvolution based on the generalized expectation-maximization method [66]. This work highlights the importance of selecting the proper laser pulse waveform model, since the maximum-likelihood solution for the range is based on the pulse shape. An investigation by Cain into the ranging accuracy limits of flash lidar systems using a negative paraboloid waveform model found an achievable range accuracy of around 6 cm, having developed an original iterative range-gain-bias joint estimator to achieve this result [14]. With the express purpose of improving the estimate of the laser pulse shape, the NOrmalized Variable Shape (NOVAS) algorithm has been developed, also by Cain and Deas [13] [24]. Furthermore, a detailed analysis on the effect of elongated pulse width in non-resolved targets, and its negative impact on range precision, has been presented in [45].

2.2 Standard Pulse Models

Improving the waveform model, and therefore peak detection, of full-waveform lidar systems is the focus of Chauve’s work, and it is suggested that waveforms be decomposed into components using a Gaussian mixture to achieve the best estimation [16]. Jutzi and Stilla also derive an improved model of the laser pulse, and introduce the Wiener filter to better estimate the reflecting surface geometry [50]. The collaboration with Kirchhof yielded an iterative method to increase the number of 3D points associated with each surface [54]. Approaching waveform fitting from a Bayesian perspective, both Hernandez-Marin [39] and Ye [107] employ the reverse-jump Markov Chain Monte Carlo (RJMCMC) methodology to improve upon range profile estimates in lidar/lidar data.

While a symmetric Gaussian curve is commonly used in lidar models, research has shown that the actual pulse is seldom thus ideal, and the symmetry assumption can lead to errors in the range estimation [56]. Improving the pulse shape estimate motivated the research in [24], and Jordan has employed a number of techniques designed to improve the range profile estimate in flash lidar data [47]. Further efforts to suppress ranging anomalies found in scanning lidar data are made in [33], where wavelet basis functions have been implemented by an expectation-maximization algorithm.

Exploiting the shape of the entire pulse waveform can help to identify specific surfaces. Correlation-based ranging using the Levenberg-Marquardt algorithm has been shown to be highly accurate, but estimation of other pulse properties (including pulse width) is also beneficial for surface identification [49]. A relevant feature vector, including pulse waveform parameters, has been used for classification of urban scenes in [17], [62]. The pixel classification in each of these papers is performed using a decision tree ensemble classifier, and achieves an error rate of less than 3%.

2.3 Image Segmentation

A basic problem encountered in image processing is segmentation, the process of partitioning an image into some non-intersecting regions such that each region is homogeneous

and the union of no two adjacent regions is homogeneous [71]. Images can represent light intensity, range depth, thermal energy, nuclear magnetic resonance, etc. The current research project proposes to segment laser range images captured by the flash lidar system. While no single segmentation method is optimal for all circumstances, a family of related techniques has been used for some time to investigate lidar-related challenges. In this section, many of these methods are reviewed and discussed in the context of the present research objectives.

Chu segments lidar data using two different data sources, the range profile and photon intensity components, and also different segmentation methods, surface fitting and statistics [20]. Edge detection methods are not preferred, as they are unsuitable for dealing with laser speckle noise. The segmentation maps are then integrated to create an improved segmentation map. A follow-up paper advances the integration algorithms of the segmentation maps [19].

Jain has also had success using range and intensity data for segmentation; however, range-based intensity histograms form the basis of the image segmentation, with an eye for potential target identification [44]. Jain is conveniently dealing with perfectly registered range and intensity images, and is integrating the two datasets using a histogram methodology. Discrimination between man-made objects and vegetation is also the subject of [73], where segmentation is achieved using texture features - the second derivative and slope - calculated directly from the range profile data.

Many image segmentation routines use training data in order to “learn” about the classes in the data set before the routine is applied to experimental data. When the upper bound, but not the exact number, of the number of classes is known, the stochastic expectation maximization (SEM) method of segmentation has been shown to be effective [63]; the simulation examines a binary classification of two Gaussian distributions. When the number of classes is known *a priori*, SEM is especially efficient. Quelle also performs unsupervised segmentation of images using SEM, but does not assume a Gaussian distribution of the data. Instead, the Pearson system allows a better fit to the histogram of the data fields and is shown to improve segmentation performance [76].

Further progress toward the completely data-driven, unsupervised segmentation method is made by Kervrann, whose algorithm is grounded in global Bayesian estimation along with Markov random field models [52]. Ten texture features describing various spatial statistics were selected for the segmentation method, which compare favorably with contemporary techniques. An unsupervised segmentation based on the skewness of the range profile, in which the ground is segmented from other objects, has been successfully applied by Bartels *et al.* to ladar data point clouds in [5, 7]; a threshold-free variant of the same algorithm is presented in [6]. Bao *et al.* have further extended this work by incorporating skewness of photon intensity [108] as well as kurtosis [4]. In [53], unsupervised segmentation is accomplished via edge detection, although it is required that each class has strong within-class invariance, as well as strong between-class separation.

Investigations into the accurate creation of Digital Surface Models (DSM) have been the motivation for several segmentation algorithms. These algorithms tend to segment LADAR data according to natural objects (e.g. trees, vegetation, or bare earth) and man-made objects (e.g. buildings, bridges, roads) for the purposes of surface fitting. While binary segmentation is popular, some methods use three or more classes. Clode uses both intensity and range profile data to segment LIDAR images into road and non-road groups [21, 22]. Samadzadegan has extended this work by incorporating classifier fusion [84]. Song achieves intensity-based four-class classification of an urban scene, due to the strong separability of the relative class magnitudes of reflectance [91]. Liu implements accurate segmentation by using spectral histograms of local windows [57].

After the regions or objects in the image have been segmented, labeling of each region is commonly done through a label classification routine. Traditional per-pixel classification often leads to “salt and pepper” effects in classification maps, although approaches such as per-pixel, sub-pixel, per-field, contextual-based, knowledge-based, and a combination of multiple classifiers are also in common use [58].

2.4 Range Data Segmentation

Much of the current work being done to improve autonomous navigation of robots and vehicles involves incorporating range data from laser radar systems. The mobile platform must be able to observe its environment, convert the data into usable features, and continue on its designated route while avoiding obstacles. Interpretation of the ladar-generated range data has been a fruitful area of research over the past several years, leading a diverse spectrum of segmentation approaches. As in this research, which extracts planar surfaces from the range data, both edge-based and region-based methods have been applied with success; however, this section focuses on work done using region-based methodology.

An early paper by Jain and Hoffman [40] segments a range image into natural object faces using a square-error criteria based on surface points and surface normals. This method provides information about the underlying object structure by avoiding using the jump edges between regions. The planar face classification is based on surface characteristics including planarity, curvature values, and eigenvalue analysis, and uses nonparametric rank tests to test for differences between the segmented patches.

More recently, Hegde’s clever approach in [38] represents range data as a tri-band color image, whereby the x, y components of each planar surface normal \vec{n} and the range depth Y are mapped to a RGB value between $[0, 255]$. From this “Enhanced Range Image” the planar segments are identified and extracted. Noise-induced range measurement errors stemming from the operation of the SwissRanger SR-3000 ladar camera operation are observed, a similar situation to the research at hand.

A similar result is obtained by Sok in [90], except that the approach is completely different from Hegde. In this paper, principal component analysis (PCA) is used to extract the planar segments are the 3D point clouds, after the point cloud is registered with a panoramic RGB images with the 3D ladar data. The range segmentation is initialized using regions of uniform color in the registered RGB image. This detail, however, is not well-suited to military and defense-oriented schemes, since color-based segmentations could potentially be defeated by camouflaging the vulnerable target.

In some autonomous navigation applications, range segmentation is more general. In obstacle avoidance scenarios, it is enough to determine whether the mobile platform can proceed in its path. Local point statistics of both range and intensity data is used in [37] and [101] to extract ground terrain information. In similar work, [25] uses range and other sensor values to make distinctions between “common” and “interesting” scenes viewed by an autonomous vehicle.

Drawing on the reflectance model proposed by Nitzan [70], Hancock proposes a laser reflectance model to detect and avoid obstacles in automated highway applications [36]. Reflectance is measured against both incidence angle and target range, and mixed results were obtained. Recent work by Gross and Jutzi advances this discussion by normalizing reflected intensity by the incidence angle of full-waveform lidar data [35]. In this paper, however, the segmentation of the image is not a priority, and neither is a discussion of the classification of the images. The data for the research is collected using an airborne laser scanner, which makes applying similar techniques to a flash lidar system a promising venture.

2.5 Classification Techniques

Pattern recognition and classification techniques assign labels (or classes) to collections of objects, or in the case of images, pixels. This can be done through unsupervised classification (where the class structure and true labels are unknown) and supervised classification (where the existing structure is known *a priori*). Supervised classification requires a training set with pre-determined labels on the objects; the classification model is built using these observations, and then applied to the unknown data, or the test set.

Observations of the objects can be collected in a measurement vector, and these parameters are used to make inferences about the class to which the object truly belongs. For example, aerial lidar data was classified into 4 classes using a Bayesian supervised learning method in [15]. Classification approaches for autonomous navigation have also been proposed by [37].

As a measure of the performance of a classification or segmentation algorithm, few tools

are more effective or more useful than the receiver operating characteristic (ROC) curve. In use since the early days of radar engineering, the ROC curve can be adapted to measure the sensitivity of the hit rate (correct classification) and false alarm rate (incorrect classification), and thus the relation between them. The simplest construction of an ROC curve involves two classes and a decision engine that classifies objects in the test set as belonging to one class or the other. The classifications are commonly displayed in a confusion matrix, the contingency table showing the differences between the true and predicted classes in the test set [11]. Ratios used to build the confusion matrix include the true positive and true negative rate, as well as false positive (false alarm) and false negative rates.

It is often necessary to compare the value of each statistic as a classifier. When it is necessary to compare the performance of different ROC curves, the area under the ROC curve (AUC) is often used, especially when a specific operating threshold is neither known nor required [11], [109]. The AUC, being a portion of the unit square, will always be within 0 and 1.0, where an AUC of 1.0 would represent a perfect classifier. The area below the random-guessing diagonal is 0.5, so a useful classifier will have an AUC greater than 0.5. The single number AUC evaluation that can then be used to compare the statistical classifiers. A notable drawback of the AUC, however, is its ignorance of the actual shape of the ROC curve in question. When two or more ROC curves have different shapes but contain the same AUC, this comparative method will score them as equal, whereas in truth we would expect one ROC curve to have differing operating points (e.g. performance), in addition to possibly different optimal points. A discussion in [11] suggests using the Neyman-Pearson criterion to maximize the specificity for the required sensitivity level. Cortes [23] has shown the calculation of confidence intervals for the AUC, while Shapiro [89] has derived the upper and lower AUC bounds in a binary hypothesis test.

The use of the ROC curve may be expanded to include multiple classes, which brings added complexity to both the calculation of the ROC curve itself and to the curve analysis. With n classes the confusion matrix becomes an $n \times n$ matrix containing the n correct classifications (the major diagonal entries) and $n^2 - n$ possible errors (the off-diagonal entries) [27]. Sampat recently presented a thorough study of three-class ROC analysis, as

well as recommending strategies for evaluation [85].

In more recent developments, Schubert [87] has expanded the use ROC curves into of multiple-class problems, culminating in the discussion of the correct classification manifold. The calculation of volume under the ROC surface (VUS), if it exists, captures only the probability of correct classification. For k classes, the volume under the ROC hypersurface will vary from $\frac{1}{k!}$ to 1, taking the value $\frac{1}{k!}$ for a completely uninformative marker (i.e. chance line) and the value 1 when the k populations are perfectly separated. Similarly, [67] extends ROC methodology to three-class diagnostic problems. The ROC surface in the three-class case is defined for diagnostic markers resulting in continuous measurements as a direct generalization of the two-sample ROC curve to three-group classification problems.

2.6 Nonparametric Statistics

Nonparametric, or distribution-free, statistics provides powerful methods for analyzing data when the underlying distribution is unknown or not defined. Although sometimes less powerful than parametric methods, nonparametric statistics have been used with success in many ranging and image applications which are pertinent to the present research. The technique of kernel density estimation (KDE), which is described in Chapter IV, is employed throughout the literature on various ladar datasets [25] [106]

Statistics of laser images are also useful in segmentation applications [4], [7], [11], [20]. Macedo has presented a statistical analysis of range data for use in obstacle avoidance situations, calculating the local range variance and skewness in order to discriminate obstacles (i.e. rocks) from traversable grass [60]. In an effort to better characterize the range distributions of forested environments, and thus improve relevant computer models, Huang has added support to the piece-wise smooth modeling technique commonly used for digital surface models [42].

Image segmentation and range profile estimation by the expectation-maximization method is again successfully accomplished by the work in [96]. Several segmentation methods use local statistics of range and intensity to compute object textures [4–6], [52], [60], [63]. Rank tests are used to test planarity of range-data segments in [40].

KDE is a well-known method for estimating the probability density function of a dataset [55]. By distributing the weight of a sample over a kernel (a known distribution, such as a Gaussian curve), KDE generates a density function that is an appealing alternative to a discrete mass function or histogram. The method can be applied to nearly any dataset, and has seen use in several papers [26] [106]. KDE is used frequently in medical imaging and statistics communities at large, but has seen relatively little use within more traditional engineering.

2.7 Lower Bounds on Estimator Performance

Parameter estimation in the context of lidar signal processing focuses mainly on the various waveform parameters, but in particular the target range, pulse amplitude, and system noise bias. Obtaining bounds on the performance of these estimators has long been a regular feature of holistic system analysis, and not unique to lidar systems. Generally, the lower bound on the estimator is found. For many precision and error-analysis applications, the Cramér-Rao bound (CRB) is the bound of choice, due in large part to its ease of computation, and good performance in the asymptotic region (high signal-to-noise ratio (SNR)) [80]. The CRB on lidar range precision was recently used to characterize lidar systems in both [45] and [65]. For a detailed derivation of the CRB, we refer the reader to [51] or [98]; a short derivation is given in Appendix A.

However, outside the asymptotic region the CRB becomes loose, and gives an overly optimistic bound on the estimator performance. For this reason alternative bounds have been introduced. In addition to the CRB, classic deterministic bounds include the Bhattacharyya, the Chapman-Robbins bound, the Barankin Bound BRB, and the Abel Bound. Two of these, the CRB and Bhattacharyya, account for the small estimation errors; naturally, their performance is good in high SNR regions, but breaks down as the SNR decreases. Conversely, the Chapman-Robbins and Barankin bounds account for large estimation errors, which accounts for performance breakdown of the estimator [80].

The performance of the estimators can also be modeled using Bayesian techniques. The Bayesian Cramér-Rao Bound (BCRB) has been put forth by [100], and also closely

predicts the performance of the estimator in the asymptotic region. More general than the BCRB, but also exhibiting tighter performance, is the Bayesian Bhattacharyya bound, which includes higher-order terms of the log likelihood function [98]. The Bayesian analog to the Barankin bound is the Bobrovksy-Zakai Bound (BZB) [9].

Finally, the final type of Bayesian bounds are the mixed or combined bounds, which integrate two or more Bayesian bounds into a single closed form. The BZB and Barankin bounds were combined in a simultaneous estimation problem, and the result is known as the Reuven-Messer bound (RMB) [43]. Recursive forms of the BCRB and BZB have been also derived and applied to a non-linear estimation problem in [77]. The Bayesian Abel bound (BAB), which combines the Bayesian Bhattacharyya constraints and the RMB constraints in a multi-dimensional optimization has recently been put forth by [80], [81]. However, the multi-dimensional problem being prohibitively complex, a single-order reduction has also been proposed. The single-order BAB reduces to a combination of the BCRB and BZB, which are both more commonly available and computable.

III. Experimental Data Collection and Waveform Analysis

The methodology for the data collected throughout the research project is described in this chapter. A description of the ladar camera hardware, and the associated parameters of the system, is given in Sec. 3.1. Sec. 3.2 discusses the layout of the field targets used in data collection. Pertinent data-preprocessing information is discussed in Sec. 3.3, including the introduction of the necessary waveform models and range data estimation. The chapter concludes with a brief summary.

3.1 Ladar Camera Properties

The ASC Portable Camera is a commercially available laser radar camera obtained in 2007 by the US Air Force Test Pilot School, and currently on loan to the Air Force Institute of Technology. As a “flash” ladar camera, it does not require a mechanically or electronically steered beam to raster across the scene in order to collect image data one pixel at a time. Instead, the entire scene is illuminated at once by the transmitting laser. Reflected light passes through the camera optics and onto the detector array, where the incoming photons are converted into full-waveform signals. The data from each pixel is collectively stored in memory, and each collection can be exported in the form of a data cuboid.

The camera, although originally designed to be self-contained, has been disassembled and mounted to a machined brass board, as seen in Fig. 3.1(a). The 3° field-of-view (FOV) lens on the ASC Portable Camera has a focal length F_0 of 250 mm and a circular aperture diameter of 120 mm. The lens is a fixed-focus lens, and is focused at infinity, which requires the targets to be The illuminating laser uses a 1.57 μm wavelength, and is equipped with a 3° diffuser which converts the laser output to a uniform illumination pattern across the entire FOV. The complete table of parameters used throughout the research experiment is given in Tab. 3.1.

For collecting data in the field, the mounting plate is oriented perpendicular to the equipment table, such that the camera and laser optics LOS is downrange. One computer connects via a 9-pin / serial port to the portable camera itself for controlling the camera

Table 3.1. 3D Ladar Camera Parameters.

| Parameter | Value |
|-------------------------------|--------------------|
| Detector Array (pixels) | 128×128 |
| Aperture Diameter | 120 mm |
| Wavelength (λ) | $1.57 \mu\text{m}$ |
| Focal Length (F_0) | 250 mm |
| Pulse Length (σ_w) | 2.7 ns |
| Pulse FWHM (τ_G) | 5 ns |
| Field of View (FOV) | 3° |
| Camera Sample Rate (f_s) | 434.5 MHz |
| Sample Period (T) | 2.3 ns |
| Range Resolution (τ_R) | 0.35 m |
| Frame Rate | 9 Hz |

operation and storing ladar images. The second laptop is used for rudimentary image processing in the field. The equipment setup is shown in Fig. 3.1(b).




Figure 3.1. (a) ASC Portable Ladar Camera mounted to machined brass plate, in storage box. (b) Camera deployed for collecting images

3.2 Target Data Collection

A series of ladar data sets were collected at a local park in late January and early February 2012. Outdoor collection is required, due to the long focal distance required by the camera. In these data collects, denoted as ‘6 Feb 12’, the first of two $2.44 \text{ m} \times 1.22 \text{ m}$ plywood board targets was erected approximately 155 m in front of the camera setup, with the second board 2.5 m behind the first board. A 2.44 m high wooden fence, along the edge of the park, was about 7 m behind the front plywood board. A frame rate of 9 Hz was used to capture 144 frames of data per collect. The frames are summed at the sample



Figure 3.2. Representative photo of data collection target.

level, yielding a $128 \times 128 \times 43$ data cuboid. The camera sample frequency of 434.5 MHz results in a sample resolution τ_R of 0.35 m. The transmitted pulse width σ_w was determined experimentally to be about 2.7 ns.

The initial target setup was a series of three nearly-parallel stepped targets (two boards and the fence), which are nearly-perpendicular to the ladar camera LOS. A similar reference setup can be seen the color photo in Fig. 3.2. A cartoon drawing of the field targets is also shown in Fig. 3.3. In the cartoon, a birds-eye view of the targets is shown. The orientation of the boards is not exact, since the targets were not calibrated for precise normal incidence. Throughout the experiment, Board 1 remained fixed.

In the second part of the experiment, the second board (Board 2 in Fig. 3.3) was rotated such that the surface normal was no longer parallel with the first board, but formed an incidence angle θ_i with respect to the ladar LOS. The board was rotated a total of three times, resulting in measurements at three different inclination angles. Although the ladar detector array is a 128×128 grid, the images in this section were cropped to 43×128 in order to draw attention to the targets under scrutiny.

3.3 Ladar Data Processing

This section describes both the raw data processing performed on the collected data, and the basic pre-processing algorithms that are used throughout the document. In Sec. 3.3.1,




Figure 3.3. Birds-eye view cartoon of data collection targets. Board 2 is titled with respect to the camera throughout the experiment.

the raw data processing is briefly described. Sec. 3.3.2 introduces the compound Poisson-noise model that is the basis of the waveform analysis performed in later chapters. Sec. 3.3.3 discusses the Gaussian-pulse shape waveform used to model the data. Finally, the basic range correlation technique is described in Sec. 3.3.4

3.3.1 Raw Data Processing.

The ASC Portable Camera collects complete angle-angle-time cuboids (or “frames”) of data at a user-defined frame rate. Each frame has dimensions of $128 \times 128 \times 20$ volume elements (voxels). A 20-sample range gate is the equivalent of about 7 m of range depth. In order to synthesize a longer range gate, the camera collects frames of data which are “stepped out” by 8 time samples; the first 12 samples in each successive frame overlap the trailing 12 samples of the previously collected frame. In the pre-processing step, the initial frame is kept in its entirety, and the “new” portions of each successive frame are appended to the original cuboid; the overlapping data is discarded. By repeatedly “stepping” out the frames, the total range gate capability of the camera is limited (in theory) only by system

Table 3.2. Per-pixel Bias Variance (counts)

| Mean | Median | Std Dev | Global Max | Global Min |
|------|--------|---------|------------|------------|
| 0.18 | 0.20 | 0.09 | 1.22 | 0 |

memory.

Effects of the internal noise bias of the ladar camera are also considered. Should significant internal bias exist, it would be manifested in the collected images as a fixed noise pattern. The internal bias data is collected by collecting 144 frames of $128 \times 128 \times 20$ data while covering the camera lens with an opaque cap. It is assumed that each of the 144 bias frames are statistically independent. A median filter is applied to the k th 2D slice in each frame (*i.e.* slice k in each of the bias frames) to remove the effects of dead/inoperative pixels. Then the variance is computed for each pixel across the 144 bias frames. The bias variance is calculated for several of the time slices, and results are consistent throughout the dataset. A summary table of the bias variance is given in Tab.3.2. The bias variance has a global maximum of 1.22 counts, and the majority of the pixels have a noise variance less than 0.2. Given the magnitude of the actual experimental datapoints when compared to the bias variance, it can reasonably be assumed that there is no intrinsic pattern noise to the camera that will be measureable in the data. Further calibrations or characterizations of the camera are beyond the scope of this research.

3.3.2 Compound Poisson Model.

Ladar imaging systems often have the capability of collecting multiple frames of the same scene, combining these frames to increase the signal-to-noise ratio (SNR). Goodman has shown that the detection of partially coherent light may be modeled using the negative binomial distribution [30]. If the number of observed photocounts at the detector is a random variable N , then the probability of k photocounts can be calculated by the pmf:

$$P\left(N = k | M, \frac{1}{1 + \frac{\mu_k}{M}}\right) = \frac{\Gamma(M + k)}{\Gamma(k + 1) \Gamma(M)} \left[1 + \frac{\mu_k}{M}\right]^{-M} \left[1 + \frac{M}{\mu_k}\right]^{-k}, \quad k = 0, 1, \dots \quad (3.1)$$

where the parameters of the distribution are M (the speckle parameter of coherence), μ_k (the expected number of photocounts) and $\frac{1}{1+\frac{\mu_k}{M}}$ (the binomial probability). The gamma function is defined by the relationship $\Gamma(n) = (n-1)!$, for any positive integer n . The expectation and variance of the negative binomial random variable in (3.1) are given by

$$\begin{aligned} E(N) &= \mu_k \\ VAR(N) &= \mu_k \left(1 + \frac{\mu_k}{M}\right) \end{aligned} \tag{3.2}$$

The speckle parameter M is defined on the range $[1, \infty)$, and in the limit as $M \rightarrow \infty$, the variance of the negative binomial random variable is given by

$$\lim_{M \rightarrow \infty} VAR(N) = \lim_{M \rightarrow \infty} \mu_k \left(1 + \frac{\mu_k}{M}\right) = \mu_k \tag{3.3}$$

which is the equivalent to $E(N)$. Equal mean and variance are a feature of the Poisson distribution. Goodman has also shown that the probability of a photcount event given incoherent object illumination may be modeled using a Poisson distribution [30]. Further characterization of the ladar camera optics, and the behavior of light, is beyond the scope of this research. For a more in-depth treatment of these two issues, McMahon [65, Ch. 2] is suggested as a reference. McMahon suggests that the sum of multiple ladar frames satisfies the assumption of incoherence; this is implemented in the experimental data collection and processing. Therefore, moving forward this research will assume that the ladar system operates under a sufficiently incoherent imaging process that can be modeled using the Poisson distribution; a similar assumption is made in [46].

Under this assumption, the photon-count of the ladar system over a time interval is an observation of the Poisson process, with expected value of photocounts μ_k proportional to the received signal waveform. This is now shown below. Let the time-varying mean value of a Poisson process $N(t), t \in [0, T]$ depend on a parameter θ as follows:

$$\lambda(t|\theta) = S(t, \theta) + B \tag{3.4}$$

where the signal part $S(t, \theta)$ is a known function of both t and θ , and B is a constant bias for the average noise power. Then the probability that k photon counts (“events”) are observed within the time interval $[t_1, t_2]$ [72] is given by

$$P(N = k|\theta) = \frac{\left[\int_{t_1}^{t_2} \lambda(t|\theta) dt\right]^k}{k!} e^{-\int_{t_1}^{t_2} \lambda(t|\theta) dt} \quad (3.5)$$

Note that for $k = 0$, (3.5) reduces to

$$P(N = 0|\theta) = \exp\left[-\int_{t_1}^{t_2} \lambda(t|\theta) dt\right] \quad (3.6)$$

The distribution in (3.5) has several names, including Poisson impulse process [30], mixed Poisson distribution [32], and the compound Poisson distribution [72].

The compound Poisson distribution is immediately applicable to the laser radar detector array. At each pixel in the detector, the observed photocount at each sample time t_i is a Poisson random variable N described by

$$P(N = k) = \frac{\left[\int_{t_i}^{t_{i+1}} \lambda(t) dt\right]^k}{k!} e^{-\int_{t_i}^{t_{i+1}} \lambda(t) dt} \quad (3.7)$$

$$k = 0, 1, 2, \dots$$

and the Gaussian pulse model is the time-varying parameter $\lambda(t)$, *i.e.* the shape of the reflected ladar waveform. Dependence on the parameter θ is implicit and suppressed.

3.3.3 Temporal Pulse Model.

The ladar pulse model proposed in this chapter is a Gaussian temporal shape, as described by [82]. Easy to model, well-understood, and with ideal properties, the Gaussian temporal pulse shape lends itself to quick analysis, since many worked solutions for Gaussian problems can be easily applied across disciplines. Throughout this research, the received (and transmitted) ladar pulse is modeled as a time-varying signal in terms of range to the

target R :

$$\lambda(t) = G_{eff} \exp \left(\frac{-1}{2\sigma_R^2} \left(t - \frac{2R}{c} \right)^2 \right) + B. \quad (3.8)$$

where G_{eff} is the effective gain parameter, σ_R is the pulse length parameter, t is the time variable, and c is the speed of light. The signal bias is denoted B . Although in this example the signal bias is a constant, it may also be modeled as a random variable. An example Gaussian pulse model is shown in Fig. 3.4. The sample pulse is a time-varying signal reflected from a target at $R = 5$ m. All other noise and distorting effects are suppressed. The Gaussian pulse width, σ_R , is related to the full-width at half-maximum (FWHM) value, denoted τ_G by:

$$FWHM_{Gauss} = \tau_G = \sigma_R \sqrt{8 \log 2}. \quad (3.9)$$




Figure 3.4. Sample Gaussian pulse model from Eq. (3.8).

The example waveform in Fig. 3.4 implicitly assumes reflection from a flat, diffuse target. In addition, it is assumed that the ladar camera line-of-sight is parallel to the surface normal of the target. These assumptions are also made in the pulse models described in [14] and [82].

3.3.4 Range Cross-Correlation.

Having established a waveform model, it remains to estimate the range to the reflecting target surface. Obtaining accurate ranging estimates remains one of the most important aspects of lidar research. However, as the present research project is oriented toward lidar image processing, it will be sufficient to use an established ranging technique, without trying to make significant improvements to it. The cross-correlation technique proposed in [82], and further explored in [13], [24], estimates both the target range and pulse width in single-echo lidar waveforms. Estimation of the pulse width will be declined initially, and the emphasis will be placed on estimating the range. Range estimates can be obtained using both parametric and nonparametric tests, and so both options are discussed.

The parametric cross-correlation algorithm of [13] computes the Pearson-product correlation coefficient for different pairs of range and pulse width, and the combination that yields the highest overall correlation value ρ is selected for the joint estimate:

$$\hat{\rho}(\hat{R}, \hat{\sigma}) = \sum_{k=1}^K \frac{(X_k - \mu_X)(Y_k - \mu_Y)}{\sigma_X \sigma_Y} \quad (3.10)$$

where \hat{R} and $\hat{\sigma}$ are the final range and pulse width estimates, respectively, X_k is the k^{th} measured waveform sample, k being an integer, Y_k is the waveform model, σ_X is the square root of the total power in the waveform samples, σ_Y is the square root of the total power in the waveform model samples, and K is the total number of samples in both the measured and modeled waveform. The sample mean of X and Y is denoted by μ . If the pulse width is assumed to be constant, then the dimensionality of (5.8) is reduced to unity.

While the Pearson correlation coefficient is an efficient, powerful method for estimating parameter values, it depends heavily on the assumption that the proposed waveform model $Y(t)$ accurately models the data, and that the noise model is additive white Gaussian noise (AWGN). If the model is deficient in some way, the calculated coefficient may lead to estimates with higher error. One possible countermeasure is to use a correlation method with less reliance on the waveform model. This would be especially useful if the observed waveform was severely undersampled, which may lead to a poor range estimate. For this

reason, a non-parametric (or distribution-free) correlation technique may be considered.

Consider a ladar waveform observation $X(t)$ that is sampled K times, resulting in a sequence of samples $\{x(t_k)\}$. Let $Y(t, \mathbf{A})$ be the waveform pulse model parameterized by \mathbf{A} , and also sampled K times. The vector \mathbf{A} is the trial values of range R and pulse length σ ; the trial values in \mathbf{A} that maximize the correlation coefficient are chosen as the estimates $(\hat{R}, \hat{\sigma})$ for the observed waveform.

Using a sample size of K , let R_1, \dots, R_K denote the ranks of the observed waveform $X(t)$ and S_1, \dots, S_K denote the ranks of the pulse model $Y(t, \mathbf{A})$. Corresponding to the parametric Pearson-product coefficient of correlation, the Spearman coefficient of rank correlation is defined as [55]

$$\hat{\rho}(\hat{R}, \hat{\sigma}) = \frac{\sum_{i=1}^K (R_i - \bar{R})(S_i - \bar{S})}{\left[\sum_{i=1}^K (R_i - \bar{R})^2 \times \sum_{i=1}^K (S_i - \bar{S})^2 \right]^{1/2}} \quad (3.11)$$

The Spearman correlation coefficient, like the parametric Pearson coefficient, ranges between -1 and 1 . If all the ranks of X and Y are equal, then $\hat{\rho} = 1$. When (3.11) is implemented for a large data sample (such as a $128 \times 128 \times 43$ ladar data cuboid), depending on the coarseness of the parameter estimate search space, the computational time may be many times greater than that of the Pearson coefficient calculations. An example follows.

Let the received waveform model $\lambda_{recv}(t_k)$ have a Gaussian pulse shape, indexed by time sample k :

$$\lambda(t_k) = \frac{G_{eff}}{\sqrt{2\pi}\sigma_R} \exp\left(\frac{-1}{2\sigma_R^2}(t_k - t_R)^2\right) + B. \quad (3.12)$$

where the pulse has some effective amplitude G_{eff} , a pulse length of σ_R , and a time-to-target of t_R ns. The pulse model in (3.12) is parameterized by the time-to-target t_R ; for simplicity, the pulse length is assumed known. In Fig. 3.6(a), the noise-free input signal is shown, while in Fig. 3.6(b) it is in the presence of AWGN. The reference waveform is shown for comparison. The model and reference waveforms plots in Fig. 3.6(a) using the signal values in Tab. 3.3. Calculating the Pearson coefficient (5.8) and the Spearman rank

coefficient (3.11) with respect to the parameter t_R , shown in Fig. 3.6(a), the equivalence of the result is clearly seen. The p -values of each coefficient are also calculated, showing that $\hat{\rho}$ is significantly different than zero at almost every t_R . For each method of correlation, the value of t_R that maximizes ρ is selected as the true value; in the present example, both the Spearman and Pearson coefficient leads to the same value of t_R . However, if the SNR is high, then the parametric correlation is much better than the nonparametric version.




Figure 3.5. Matched filtering comparison of parametric and non-parametric method.

This example introduces the concept of nonparametric correlation for future use with ladar pulse waveform fitting. A notable drawback to using the nonparametric correlation is the increased processing time as compared to the parametric correlation method. Processing a complete ladar dataset using the nonparametric correlation method took nearly 6 times as long (181 sec) as using the parametric correlation (31.5 sec). Additional analysis also shows that using the parametric correlation method for range estimates yields satisfactory results.




Figure 3.6. Correlation coefficients and p-values for matched filter output.

Table 3.3. Waveform parameters for the sample plots of Fig. 3.6 and Fig. 3.6.

| Parameter | Value |
|------------|-------|
| σ_s | 5 |
| σ_r | 2 |
| G_s | 5 |
| G_r | 5 |
| B_s | 5 |
| B_r | 3 |
| R_s | 2.5 |
| R_r | 2.5 |

However, using the nonparametric rank correlation for applications involving undersampled signals, or when the waveform model is poorly developed, may be preferable to parametric methods, and could be investigated in the future.

3.4 Summary

This chapter briefly described the parameters of the ASC Portable Ladar camera that was used throughout the research, and established the preferred waveform model. In addition to describing the geometry of the target layout, the necessary data pre-processing steps were also discussed. The compound Poisson distribution, with a Gaussian time-varying mean, is selected for the system model. Simple cross-correlation range estimation is introduced, and is discussed in both parametric and nonparametric versions.

IV. Nonparametric Ladar Image Segmentation

What does a laser radar image consist of? As is the case in many engineering applications, the only honest answer is: it depends. The first generations of ladar systems collected point cloud data, consisting of ordered pairs of peak amplitude (photo-count) and time-to-target (or, equivalently, target range). As it was (and remains) common for ladar systems to be used for ranging applications, the $2D$ “image” formed by the data is a peak, or principal, range image; an example is shown in Fig. 4.1. Visualizing the $2D$ pulse amplitude “image” follows naturally from the range data. Since these images are obtained with a minimal amount of signal processing, and are also calculated at each individual pixel, they can be referred to as “first-phase” images.

In contrast with earlier technology, contemporary full-waveform ladar systems (including the ASC Portable Camera) provide the opportunity to extract additional range information from the illuminated scene. Whereas the point cloud data consists of a single peak range value, the full-waveform data may contain multiple pulse echoes, indicating surface diversity in the line of sight (LOS). Pulse shape distortion commonly occurs when the ladar pulse reflects from angled or rough surfaces, and is especially associated with reflection from vegetative targets, including tree foliage. Measurement of the pulse shape distortion has been shown in [48], although this was done in a tightly controlled laboratory environment using a pencil-beam laser to illuminate a plate target with ideal backscatter characteristics.




Figure 4.1. Sample “first-phase” principal range image.

Considering the flash ladar camera used in this research, a single dataset consists of an angle-angle-time cuboid, not merely point cloud data. A flattened, $2D$ intensity image,

such as the example in Fig. 4.2, is formed by the summation of the cuboid data in the time dimension. This visualization necessarily obscures any ringing or multiple echoes that might be observed by looking at the complete waveform. However, should signal analysis conclude that M multiple echoes are present in the waveform, it would be possible to visualize an M -ary range image in conjunction with the principal range image described earlier. These images, which are created through signal analysis of singular, full-waveforms at each pixel, are the “second-phase” images.




Figure 4.2. Sample “second-phase” 2D intensity image.

Finally, “third-phase” images are created by analyzing waveforms from multiple pixels. The information conveyed by the third-phase image is derived from the interaction of several pixels at once, and cannot be known through the observations at a single pixel. One example of this is a fitted-plane image. Commonly used in digital terrain mapping (DTM), an algorithm uses a local window of principal range values to find the best-fit surface across the selected points. The plane is defined locally by the normal surface vector \vec{n} , which also forms an inclination angle θ_i with respect to the illumination axis, defined by the global coordinate system unit vector \vec{e}_{unit} . Use of surface normal vectors has recently been explored in [38], where the geometric (surface normals) and coordinate data were fused into a RGB-color based Enhanced Range Image, which enhances object edges and surfaces, and is used as the basis of segmentation. Hegde’s approach, however, is entirely different than the solution discussed in this chapter.

The present solution recognizes that each pixel in the detector array, in addition to having a set of observations (i.e. photocounts) stored in memory, is associated with both a local surface and the inclination angle of the surface. Additionally, knowledge of the incli-

nation angle of each target may provide insight into the backscatter reflectance properties of the target surface. Images formed by these datapoints, the surfaces, inclination angles, and reflectance characteristics, are just further examples of third-phase images. A study related to intensity normalization, using local surface incidence angles and range data, is presented in [35], although different methodology is used throughout the paper.

Throughout this discussion of ladar data, it is clear that there is no single “image” that captures all the available information about the laser-illuminated scene. Any discussion about performing image processing on a ladar dataset must recognize this peculiar aspect of the data. Having a defense-oriented objective in mind, isolating *targets of interest* within the illuminated scene will be the most appropriate objective. Again the question arises: what is *interesting*? A recent paper described one attempt to differentiate “common” from “interesting” scenes [25]. In a military framework, an interesting object usually meets one or more of the following criteria: it is different from everything around it; it is hidden or obscured in some way; and/or it has anomalous characteristics.

Keeping these things in mind, the image processing techniques presented in this chapter will be shown to be valuable to military remote-sensing applications. The chapter opens in Sec. 4.1 by describing in full the nonparametric image segmentation routine that was developed for this research. The method is based on the range data of the ladar image, from which a probability density is calculated. Object points are derived from the range density, and the segmentation is accomplished using familiar image processing techniques. Additional third-phase images are extracted during the segmentation process, and the usefulness of the data is examined. In Sec. 4.2, the data is segmented according to a more traditional region growing technique, and the results are compared with the methods of Sec. 4.1. Generation of third-phase images is described in Sec. 4.3. The chapter concludes in with a summary in Sec. 4.4.

4.1 Nonparametric Range Data Segmentation

It is a truism to state that an object is located in the space that it occupies. However, any effort to identify (or segment) a target of interest that may be present in a ladar dataset

must recognize that each object in the image occupies a definite location in space, that is, at a specific range distance. Therefore, the segmentation routine presented in this section will be initially derived from the principal (highest peak amplitude) range image extracted from the ladar data cuboid.

Some goals to keep in mind with respect to image segmentation are: every pixel should be segmented into exactly one region, and no pixels should be unsegmented. Also, segmented regions should be contiguous as well as distinct from each other.

4.1.1 Kernel Density Estimation.

Consider that every pixel in the ladar data cuboid has a principal range associated with the detected waveform. A waveform lacking a principal range would correspond to a “no echo” return, or a pulse reflection beyond the range gate; however, we will decline that special case at the present time, and assume that every pixel detects a target within the overall range gate. A histogram of the principal range values at each pixel in the range image might look something like Fig. 4.3(a). However, if the number of bins is reduced, then the histogram may look more like Fig. 4.3(b). Visualizing the empirical density of the range in this way is primitive, and is wholly dependent on the choices of bin widths and number of bins.




Figure 4.3. Sample range density histograms

Kernel density estimation (KDE) spreads out the weight of a single observation in a

plot of the empirical density function [55]. This nonparametric technique for constructing probability density functions has been the subject of many inquiries, and can be applied to nearly any collection of data, regardless of the source [2] [25] [10]. In Figs. 4.3, the probability mass of each observation is spread uniformly throughout its appropriate bin. By using a different kernel, the probability mass of each item is distributed in an alternative, usually symmetric, way. The density function calculations are straightforward. Let the density estimate of a sample \mathbf{X} be given by

$$\hat{p}(x) = \frac{1}{nh_n} \sum_{i=1}^n K\left(\frac{x - x_i}{h_n}\right), \quad i = 1, \dots, n \quad (4.1)$$

for $X_i = x_i$, and where n is the total number of observations. The probability mass of each x_i is assigned according to the kernel function K , and the smoothing function h_n represents the bandwidth. The kernel function of the basic histogram is just a constant, and the density estimate is discrete. By choosing a different kernel, for example the normal (or Gaussian) kernel, the density estimate is smooth and continuous; this outcome is shown in Fig. 4.4(b). Interestingly, for the example data output from the ASC Portable Ladar Camera, the optimal smoothing bandwidth (in a Gaussian sense) of 0.35 m is also the length of the range resolution, $\tau = 0.35$ m. A smaller bandwidth, such as half of the range resolution, or $h_n = \tau/2$, reveals hidden modes in the data. By increasing the bandwidth to twice the range resolution, or $h_n = 2\tau$, the minor features in the density estimate are obscured. Each of these estimates are shown in Fig. 4.4(a),(c). Constructing the density estimate via non-parametric (or distribution-free) statistics is appropriate, since there is no standard probability distribution that would accurately describe the range density. The next section details how the segmentation is performed using the bins of the range density.

4.1.2 Density Range Bins.

The range density estimate that is constructed using the Gaussian KDE describes the probability mass of objects located at any given range. In Fig. 4.5, the peaks and valleys




Figure 4.4. Sample range density using Gaussian kernel. Bandwidth values: (a) 0.17 m (b) 0.35 m (c) 0.70 m

of the density estimate are indicated by red triangles and black diamonds, respectively. By including the endpoints of the density estimate as minimums, the chart shows a series of 7 peaks (or local maximums) 8 valleys (or local minimums). A difference threshold of 0.005 is used to identify the local maximum. The bandwidth selection of 0.17 m appeared to identify hidden modes, so this shorter distance was selected over the “optimal” value of 0.35 m.

To create the bins, the density function is selected piecewise between the local minimums. For example, in Fig. 4.5, the first “bin” begins at the the minimum x value, 152 m, and extends to the first minimum at 157.1 m. The rest of the density estimate can be parsed in the same fashion, resulting in 7 bins of unequal length, as measured between the local




Figure 4.5. Sample range density estimate, calculated using Gaussian kernel. Bandwidth: 0.17 m

minimums. Viewed in this way, the complete density function is composed of piecewise density estimates, with each bin corresponding to a region of range probability. The range data from the original image corresponds at the pixel level to a level of probability in the density estimate. This establishes the empirical likelihood of a particular pixel being associated with a specific principal range. Exploiting the density estimate in this way has not, according to a contemporary search of the literature, been attempted in the context of lidar range data.

4.1.3 Boundary Tracing and Segmentation.

Each range bin from the density function, where up to n total bins are partitioned between the local minimums, corresponds to a set of points in the principal range image. The points are defined by their (x, y) pixel location and their z -value range. The pixels that correspond to the probability in each of the n bins in the estimate are selected one at a time for processing. For example, viewed as a 3D surface image in Fig. 4.6, it is clear that the nearest object is in a distinct range bin by itself. The portion of the principal range image that falls into the $n = 1$ bin is shown in Fig. 4.7(a). In general, each “slice” is

composed of regions of both empty space (outside-the-slice) and in-slice data. Basic image morphology is used to remove small (< 10 pixels) objects from the image. The in-slice data that remains is organized by into m contiguous regions $\mathcal{R}_n = \{r_1, r_2, \dots, r_m\}$, where r_k is the k th contiguous region in the n th slice, and $m \geq 1$. In the example slice of Fig. 4.7(a), there is a single region, so $m = 1$. The total number of segmented regions in the image is the sum total of the elements in every \mathcal{R}_n , which shall be defined as the integer R_Σ .

From here, the boundary of each region r_k is extracted using the MATLAB[®] function `bwboundaries`, which itself implements a tracing algorithm found in [29]. For the current example, the binary mask and the associated boundary are shown in Fig. 4.7(b). The set of boundaries for each slice is denoted $\mathcal{B}_n = \{b_1, b_2, \dots, b_m\}$, since each region r_k has an associated boundary b_k .

The boundary tracing is continued for all of the m regions in each of the n range density slices. The complete set of range data boundaries is denoted $\mathfrak{B} = \{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_n\}$. The composite boundary set \mathfrak{B} is overlaid on the second-phase intensity image in Fig. 4.8. Since each closed boundary traces the border of a continuous region rather than the border between adjacent regions, a “double-line” effect is observed. However, the lines that define the boundary are for the most part defined on adjacent pixels, so there is no true gap between the contiguous regions.

Fig. 4.8 is the complete range data segmentation of the ladar cuboid using the $\tau/4 = 0.17$ m bandwidth. In the image, the segmented objects are the two board targets, the foreground (grassy field), the background (tree branches) and the fence. The fenceline across the image is segmented into two regions, due to the fence being an elongated target that extends across two range bins.

As a comparison, the range data segmentation can be performed using a different range density bandwidth. Referring to the range density of Fig. 4.4(b), the bandwidth is $\tau = 0.35$ m, causing the total number of bins to decrease by one. The resultant composite range data segmentation is shown in Fig. 4.9. In this segmentation, the fenceline is encapsulated into a single range bin, even though the surface is not completely planar.

In the next section, several examples of the range data segmentation method are pre-

sented and analyzed.




Figure 4.6. Sample range surface image.




Figure 4.7. $n = 1$ range density bin images.




Figure 4.8. Composite boundaries \mathfrak{B} , overlaid on 2nd phase intensity image. Density bandwidth: 0.17 m




Figure 4.9. Composite boundaries \mathfrak{B} , overlaid on 2nd phase intensity image. Density bandwidth: 0.35 m

4.1.4 Examples.

The ladar data cuboid analyzed throughout this section is the ‘6 Feb 12’ dataset, collected a local park in late January and early February 2012 using the ASC Portable camera. Further specific information on the dataset and the camera is found in Chapter III. Before the segmentation process begins, each 128×128 2D slice of data (in the time-direction) is median filtered in a 3×3 neighborhood of pixels. This filtering operation is intended to reduce the speckle noise across the image.

In the ‘6 Feb 12’ datasets, the 2.44 m \times 1.22 m plywood board target was erected approximately 155 m in front of the camera setup, with the second board 2.5 m behind the first board. A 2.44 m high wooden fence, along the edge of the park, was about 7 m behind the front plywood board. Principal range images of all the 6 Feb 12 trials are shown in Fig. 4.10. The alternate colormap of Fig. 4.10(a) is due to the data being collected using a longer range gate. Range density estimate and segmentation of the first trial is shown in Fig. 4.11 using a default bandwidth of 0.35 m. In the first trial, the two targets are basically perpendicular to the camera, so no targets are split between range bins.

For the remaining 3 trials, the second plywood board was successively angled with respect to the ladar camera. Using a standard bandwidth of 0.35 m, the range density estimate and composite range segmentation is shown in Figs. 4.12-4.14. Of note in each figure is the second peak of the range density estimate. This local maximum corresponds to the second plywood target. As the aspect angle increases, the board takes up fewer pixels in the image; thus the probability at that range diminishes. Only in Fig. 4.14(b)

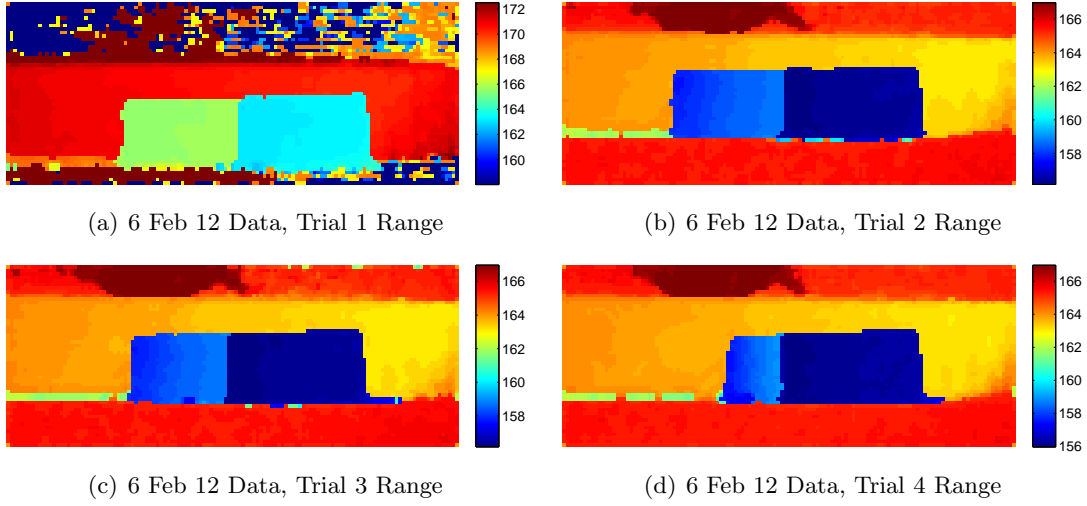


Figure 4.10. Principal range image of 6 Feb 12 range data, trials 1–4.

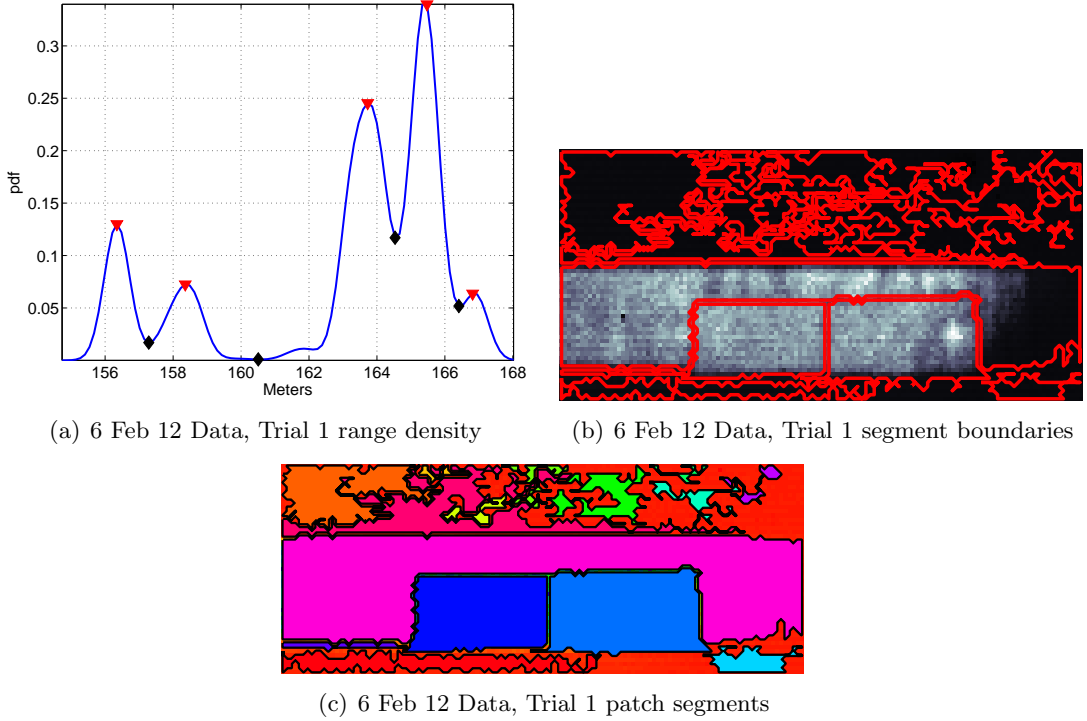


Figure 4.11. Density ($BW = 0.35$ m) and segmentation of 6 Feb 12 data, target 1.

does the probability decrease nearly to the point of being negligible. In any case, in each of the datasets, the image is completely segmented - that is, every pixel belongs to a specific region, and all boundaries are unbroken.

To observe the total effectiveness of the segmentation method, the results are compared

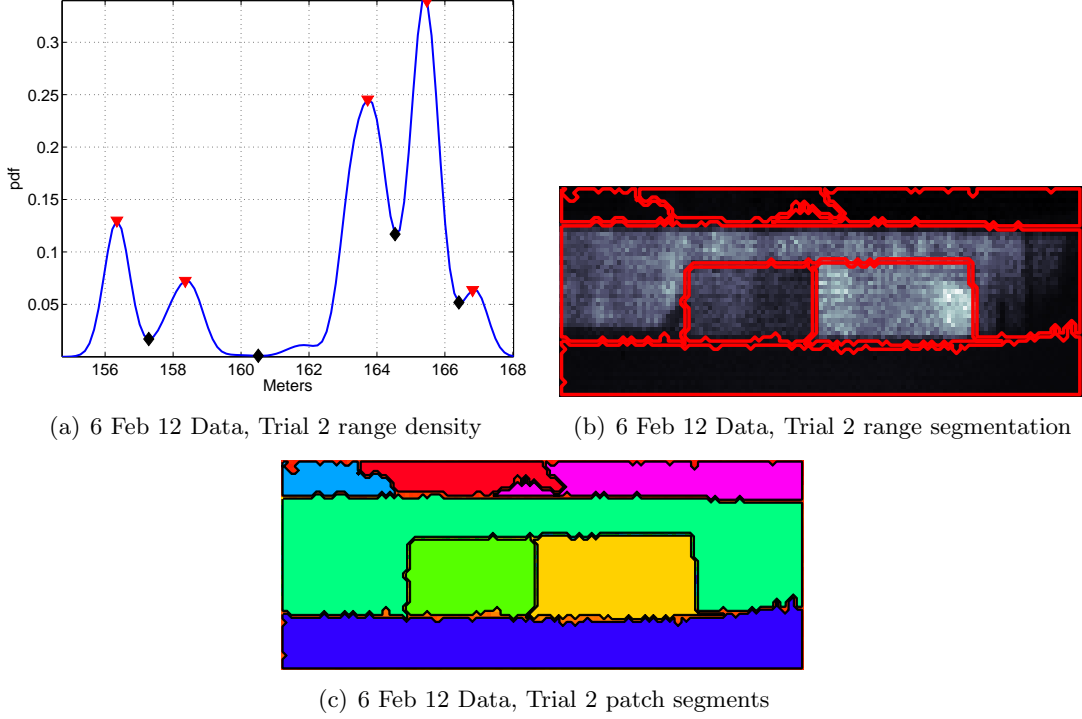


Figure 4.12. Density ($BW = 0.35$ m) and segmentation of 6 Feb 12 data, target 2.

to an edge detection filter. A Canny edge detection filter, while it has known drawbacks, provides a reasonable metric for visual comparison of segmentation results [61]. It is especially effective in identifying weak edges. The edges of the principal range data image for each of the 4 trials, generated from the corresponding subplots in Fig. 4.10, are shown in Fig. 4.15. None of the edge maps completely segment the range image with unbroken boundaries, and in the case of the Trial 1 data, the edge detection fails to distinguish between the two target boards. Use of Canny edge detection to segment the principal range image is clearly inferior to the novel KDE-based segmentation proposed in this section.

4.1.5 Summary.

In this section, nonparametric range data segmentation was accomplished on a series of datasets acquired with the ASC Portable Ladar camera. The segmentation is initialized using a Gaussian-based KDE method, the result of which is a probability function of the ladar principal range image. Bandwidth selection for the density estimate is a crucial component

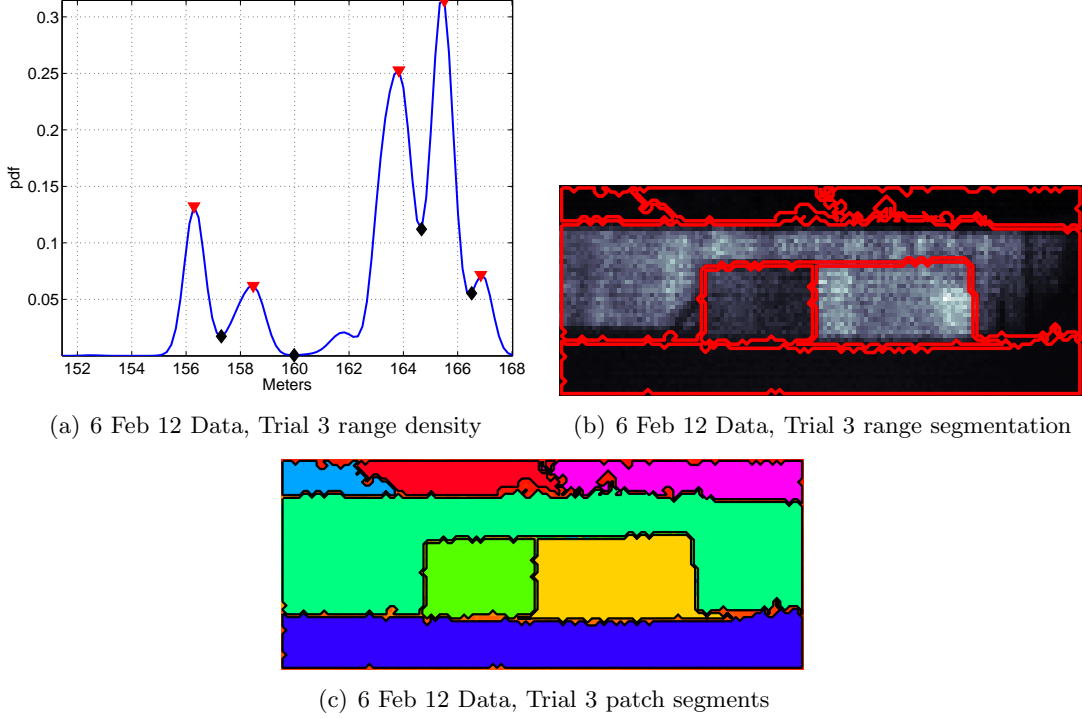


Figure 4.13. Density ($BW = 0.35$ m) and segmentation of 6 Feb 12 data, target 3.

of the segmentation process, as was shown. The lidar camera range resolution parameter τ is an appropriate choice of bandwidth to obtain acceptable segmentation results, and this was selected as the default value. Range data segmentation is accomplished successfully using nonparametric kernel density estimation.

4.2 Iterative Region Growing Segmentation

In this section, a comparative analysis is presented that illustrates the results of an iterative region growing segmentation method, based in part on recently proposed algorithms by Taylor [95] and Chen [18]. A region growing approach is offered here as an alternative to the nonparametric segmentation derived in the previous section. The seeded region growing (SRG) method is one of the most simple and popular algorithms for segmentation, so it is appropriate for use in comparing to the novel nonparametric segmentation method. A nonparametric rank-sum test is used at the conclusion of the section to compare the classification rates of each proposed segmentation methodology. The steps of the iterative




Figure 4.14. Density ($BW = 0.35$ m) and segmentation of 6 Feb 12 data, target 4.

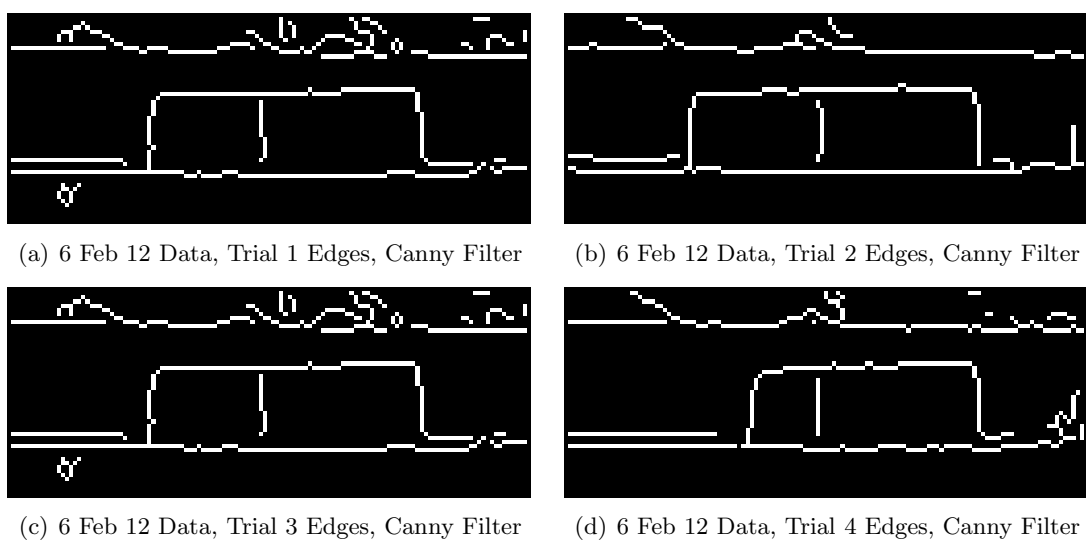


Figure 4.15. Canny edge detection of 6 Feb 12 range data, trials 1–4.

region growing algorithm are described in Sec. 4.2.1, and the results of the segmentation are discussed in Sec. 4.2.2.

4.2.1 Segmentation Algorithm.

Taylor has suggested initially segmenting range data at jump boundaries and creases, since a sudden change in range (or orientation) can indicate an abrupt change in a planar face [95]. This criteria is also well suited for segmenting man-made objects (also typically planar) in laser radar data. The initial segments will be used as seed regions, which will grow to include the unlabeled pixels in the remainder of the image. The input to the region-growing segmentation is the principal ladar range image. Calculating approximate surface normals at each pixel using a sliding window approach is a straightforward approach, but non-physical results are obtained at the jump edges of the object boundaries.

The algorithm proceeds as follows:

1. Approximate the surface normal \vec{n}_i at each image pixel using a 5×5 sliding window over the range image. PCA (Sec. 4.3.1) can be used to find the components of \vec{n}_i with respect to the ladar LOS. The incidence angle at each pixel is found from the dot product $\vec{n} \cdot \vec{e}_z = \cos \theta_i$.
2. Points satisfying $\theta_i > \phi_m$, where ϕ_m is a user-defined threshold, are classified as jump regions. These points are discarded and removed from the segment map.
3. The remaining regions represent the planar faces in the image, and boundary tracing is used to identify the contiguous regions. The surface normal in each region is recalculated using all the points in the region; this should be more accurate than the first calculation, since the jump edges are not included. Statistics of each region, including the mean range and range RMSE, are calculated using all the points in each region.
4. Using the existing regions as seeds, an iterative region growing loop assigns unlabeled range elements to the seed regions. Two stopping conditions are implemented for the algorithm, using both edge information and a range RMSE constraint. Unlabeled pixels are added to an adjacent region if the range value r_i is sufficiently close to the mean range \bar{r} of a region: $|\bar{r} - r_i| < d_m$, where d_m is a user-specified threshold scaling

of the range RMSE. Adjoining pixels are not added to the region if they coincide with the edge boundaries of the range image, which are found using the Canny edge filter.

5. The resulting regions are considered for overlap with other regions, and subjected to a visual analysis. Overlapping regions are merged manually until a satisfactory segmentation of the image is completed.

4.2.2 Results.

The iterative region-growing algorithm presented in the previous section has been applied to the '6 Feb 12' dataset, and the results of the method are presented in this section. The region seeds were initialized by setting the angle threshold $\phi_m = 60^\circ$. This choice of angle threshold ensured that the expected planar targets were not discarded as either background or highly angled regions. After the threshold was applied to the global surface normal image, 5 distinct seed regions remained, and this was the case for each of the trial datasets. The discarded pixels which were not included in the seed regions are unlabeled.

In the iterations, regions were taken one at a time. Unlabeled pixels adjacent to the region perimeter were added to the region provided the range RMSE and edge constraints were satisfied. It was necessary to manually adjust the scaling of the range RMSE constraint between datasets in order to achieve segmentation of the tilted target board. The threshold d_m was between $1.1 \times RMSE$ and $4 \times RMSE$ depending on the dataset, where the RMSE was calculated separately for each segmented region. The edge constraint was determined using the Canny edge filter of the principal range image, shown in Fig. 4.15. If a candidate pixel, after satisfying the RMSE condition, was found not to be a boundary pixel, it was finally added to the region. Region growing iterations continued until the region, bounded by edges or non-candidate pixels, grew by less than 5%. The remaining regions were taken in turn.

Background segments, when they occurred, were manually selected for merging after the completion of the region growing iterations. The planar targets imaged by the camera were all properly and correctly segmented by the algorithm. The segmentation of each trial

dataset using the iterative SRG technique is shown in Fig. 4.16, while the nonparametric segmentation from the previous section is repeated in Fig. 4.17. At a qualitative level, the segmentation results are comparable, and this is a positive result. The segmentation of the first dataset, in Fig. 4.16(a), does not cleanly identify the region boundaries as the corresponding nonparametric segmentation in Fig. 4.17(a). Similarly, the segmentation of the second (tilted) board in Fig. 4.16(d) is rough around the edges. This is due in part to the low SNR of the waveforms in this region, and the fact that the RMSE threshold d_m was highest in this region, so as to correctly include the whole of the extended target. The segmentation is in contrast to that of Fig. 4.17(d), which is the nonparametric version. Segmentations of trials 2 and 3 are similar for both segmentation algorithms, and are shown in Fig. 4.17(b) and Fig. 4.17(c), respectively.

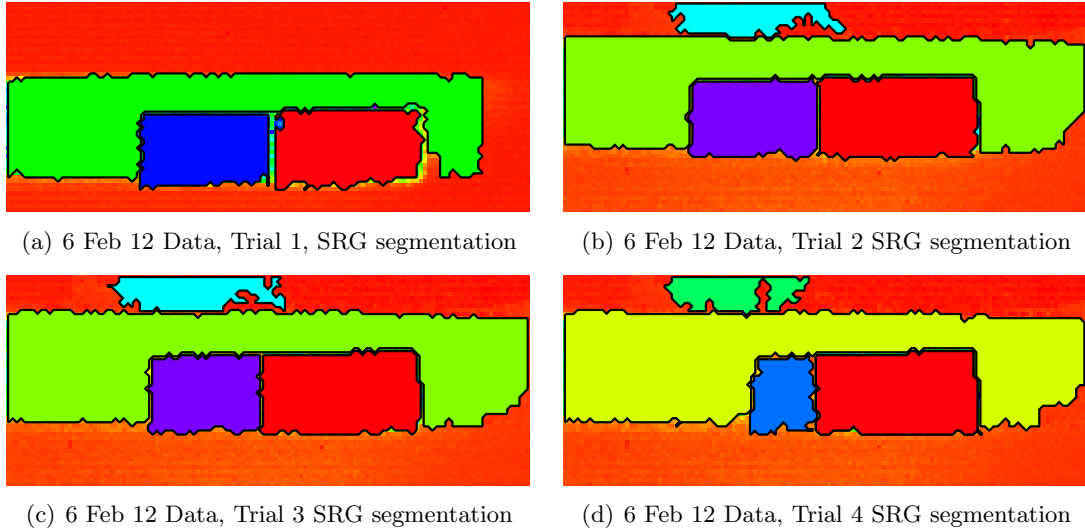


Figure 4.16. Segmentation using iterative SRG on 6 Feb 12 range data, trials 1–4.

When the segmentation algorithms are compared to a hand-labeled segmentation, the effectiveness of the methods can be numerically evaluated by classification rate (CR). The segmented regions of interest for the comparison are the fixed board, the angled board, and the fence. A manual segmentation of each region in each of the 4 trial datasets provides the benchmark for the comparison. The areas of the segmented regions that result from the nonparametric and seeded region growing algorithms are computed and compared with the manual segmentation. Since the stopping criteria of each method is different, the comparison

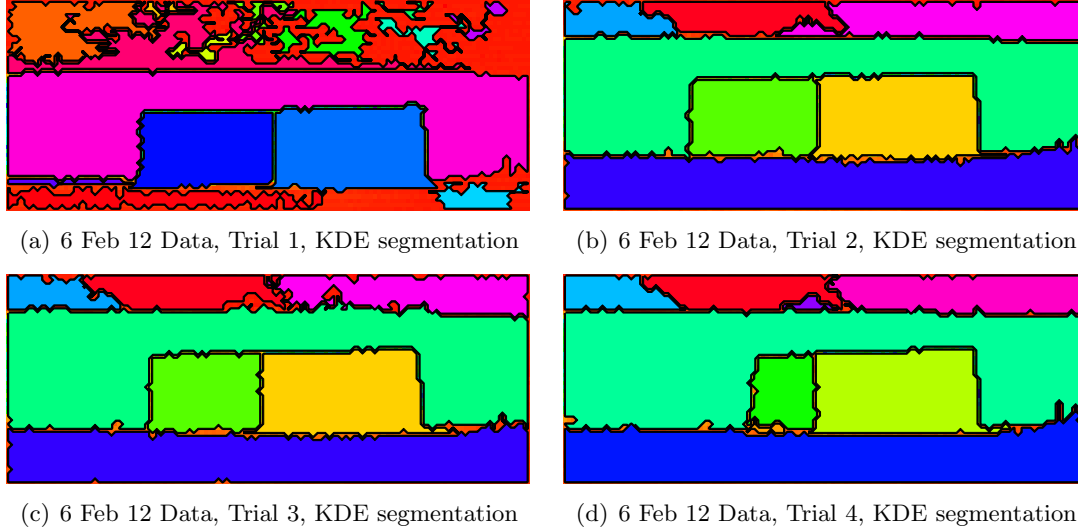


Figure 4.17. Segmentation using nonparametric range KDE, on 6 Feb 12 range data, trials 1–4.

Table 4.1. Area & Classification Rates (CRs) of Fixed Board Segmentation.

| | Trial 1 | Trial 2 | Trial 3 | Trial 4 |
|---------------|------------|-------------|-------------|-------------|
| Manual | 777 | 796 | 749 | 763 |
| Nonparametric | 747, (96%) | 796, (100%) | 811, (108%) | 820, (107%) |
| SRG | 667, (85%) | 707, (89%) | 717, (95%) | 770, (101%) |

does not determine whether the same pixels were segmented into the same regions; this outcome is assumed to generally be true, as a visual inspection of the segmentations will indicate. Rather, the comparison of area gives an idea of how restrictive or generous each segmentation method might be.

The results in Tab. 4.1 show the number of segmented pixels in each trial that correspond to the fixed board target. Using the manual segmentation as a benchmark, there is a small amount of variation across the trials, but on the whole each method is consistent. The nonparametric method routinely segments a larger area than the manual segmentation, while the SRG segment is smaller than the benchmark. Since the orientation of the fixed board did not change from trial to trial, this comparison shows the consistent performance of each segmentation method.

The CRs of the angled target board are given in Tab. 4.2. The facing area of the target board decreases with each trial, corresponding to the target rotation. As before

Table 4.2. Area & Classification Rates (CRs) of Angled Board Segmentation.

| | Trial 1 | Trial 2 | Trial 3 | Trial 4 |
|---------------|------------|------------|------------|-------------|
| Manual | 619 | 589 | 524 | 256 |
| Nonparametric | 603, (97%) | 585, (99%) | 511, (98%) | 264, (103%) |
| SRG | 520, (84%) | 545, (93%) | 469, (90%) | 254, (99%) |

Table 4.3. Area & Classification Rates (CRs) of Fence Segmentation.

| | Trial 1 | Trial 2 | Trial 3 | Trial 4 |
|---------------|--------------|--------------|--------------|--------------|
| Manual | 2120 | 2131 | 2243 | 2497 |
| Nonparametric | 2126, (100%) | 2260, (106%) | 2413, (108%) | 2554, (102%) |
| SRG | 1812, (85%) | 2078, (98%) | 2152, (96%) | 2288, (92%) |

in Tab. 4.1, the nonparametric segmentation generally results in a larger area than the manual labeling, while the SRG segments are smaller. Since the stopping criteria of the SRG method was based on both edge detection and range-RMSE, it is a positive outcome that the segmentation of the board in Trials 3 and 4 are so close to the benchmark.

Finally, the CRs of the fence segmentation are found in Tab. 4.3. The performance of each method is similar to that of the previous two target segmentations, in that the nonparametric method segments larger regions compared to the manual labeling, and the SRG method yields slightly smaller regions. In all three tables, the performance of the SRG method for Trial 1 data is consistently lower than in every other entry. This is likely due to inconsistencies with data collection, and alternative ladar camera and range gate settings.

To confirm that there exists measurable difference between the nonparametric and SRG segmentation algorithms, a rank-sum test was used to test the similarity of the CR. A rank-sum test proposes the null hypothesis that the samples in two sets come from identical distributions with equal medians, versus the alternative that the medians are unequal. For this test, each set of data will be the CRs of each of the segmentation methods: nonparametric (\mathbf{X}), and seeded region growing (\mathbf{Y}). The test is constructed as:

$$\begin{aligned}
H_0 : & \text{CRs of segmentation methods are equal} \\
H_1 : & \text{CRs of segmentation methods are not equal}
\end{aligned}
\tag{4.2}$$

The rank-sum test statistic for these two samples is 210, so the null hypothesis is rejected

at the 5% significance level. It can be concluded that the segmentation methods yield significantly different CRs when compared to a hand-segmented benchmark.

4.2.3 Summary.

Consideration of a more popular and well known segmentation approach using iterative SRG has been shown to yield comparable results to the nonparametric range-KDE method. This implies that the nonparametric segmentation approach can be considered at least as good as the more traditional SRG. Datapoints extracted from the two segmentation outcomes, such as planar incidence angle, are also roughly similar, and any discrepancies are within the error bounds of the measurements. One drawback of the SRG approach is the requirement for user involvement in several tasks, including choosing the surface normal angle defining a 'jump edge', the scaling of the RMSE constraint, and manually merging overlapping regions. The nonparametric approach reduces the number of user interactions without losing the goodness of the segmentation. In addition, the considerations of stopping criteria, which are common to all iterative methods, add a level an additional complexity to this segmentation approach that is undesirable for the purposes of this research.

4.3 Third-Phase Ladar Image Analysis

Extracting additional layers of data from the first- and second-phase ladar images is the objective of this section. A third-phase image has been defined as an image created by analyzing waveforms from multiple pixels. By contrast, first- and second-phase images are created using waveforms from individual pixels. Two types of third-phase images are discussed in this section: the surface normal image and the material reflectance image. Although a surface normal exists for each pixel in the image, the components of the vector can only be determined by using information (in this case, range) from adjacent pixels. Material reflectance is dependent on both the material and the illuminating wavelength, although the backscattered intensity is dependent on the aspect angle of the object with respect to the laser. The reflectance, therefore, can only be estimated using knowledge of the aspect angle, which satisfies the definition of the third phase image.

This section describes the calculation of the surface normal image through plane fitting in Sec. 4.3.1, and the material reflectance image in Sec. 4.3.2. Examples, using the 6 Feb 12 data, are provided at the end.

4.3.1 Plane Fitting.

After the segmentation has been performed on the regions in each bin, the range data of the points within each region are analyzed. The points of the contiguous regions fall, by design, within some “bin” that has been defined on the range density curve. Each point is described by its (x, y) pixel location as well as its range value in the z -axis. Fitting a plane to the points in each region requires a number of assumptions about the object in question. First, the object is assumed to be locally planar within the pre-defined range “bin.” The distinctive features of curved surfaces, multi-faceted surfaces, and irregular-shaped objects will be lost during the plane-fitting procedure. Secondly, it is assumed that the object is completely located within a single range bin. Extended targets, such as a long fence line or some other elongated object, may be divided into separate range bins during the density estimate procedure. The plane fitting then will fit separate planes to the object, which in light of these caveats, may or may not be indicative of the target’s true orientation.

These assumptions notwithstanding, the plane fitting operation is straightforward. The first step requires transforming the (x, y) pixel indices into actual distance markers. This is accomplished in the following manner. The camera lens FOV is known *a priori* to be 3° , which corresponds to a half-angle of 1.5° . Within each segmented region r_k , the median range value, denoted R_{med} , is selected. The transverse scene width of the camera FOV, denoted d_{trans} is then calculated to be

$$d_{trans} = 2R_{med} \times \tan(FOV/2), \quad (4.3)$$

which is measured in meters. As an example, if the median range of a segmented region is 160 m, and the FOV is 3° , then the transverse width of the scene is 8.34 m. Dividing the transverse scene width by the number of pixels in a single dimension of the 128×128

detector array yields the distance per pixel, a value of 0.066 m, or 6.6 cm. This distance is valid in both the x and y direction, due to the radial symmetry of the camera lens.

Per-pixel calibration can also be accomplished using the known sizes of the reference board in the image. The dimensions of the front board are 2.44×1.22 m, dimensions which correspond in the image to 36×19 pixels. Using these dimensions results in each pixel corresponding to an area of 6.8×6.4 cm, reasonably close to the estimate derived from the FOV calculation.

Using the (x, y) pixel location as an index, each point is labeled with a 3-dimensional spatial location in meters. If a region is composed of N points, then the set of data points in the region r_k is $\mathcal{D}_k = \{x_i, y_i, z_i\}$ with $i = 1, 2, \dots, N$, and has dimensions $N \times 3$. Principal component analysis is now used to calculate the surface normal for the best-fit plane of \mathcal{D}_k .

Principal component analysis (PCA) is a powerful tool used to form an orthogonal basis for the space of the data, and, by extension, find the best-fit planar surface through the collection of points. The 3D feature extraction algorithm developed by Sok applied PCA to ladar data, although the image in that research was first segmented using color information [90]. The coefficients of the first two principal components define the basis vectors $\langle \vec{u}_1, \vec{u}_2 \rangle$ of the plane, and the coefficient of the third principal component defines the normal vector \vec{n} of the plane. It is the normal vector that is of most interest to this research. The mean-squared error (MSE) of the resulting plane gives an indication of the variation in the fitting, and can be used to quantify the effectiveness of the PCA.

The diagram in Fig. 4.18 shows how the surface normal vector, \vec{n} , forms an angle θ_i with the z -axis. The incidence angle θ_i has radial symmetry around both the z -axis and the ladar boresight, according to the scenario geometry. By defining the unit vector \vec{e}_z as parallel to the z -axis, the cosine of the angle between the two vectors is given by the dot product

$$\vec{n} \cdot \vec{e}_z = \cos \theta_i \quad (4.4)$$

In a new, third-phase image, the inclination angle θ_i is mapped to the (x, y) position of the points in \mathcal{D}_k , meaning that all the pixels in a region are given the same surface normal

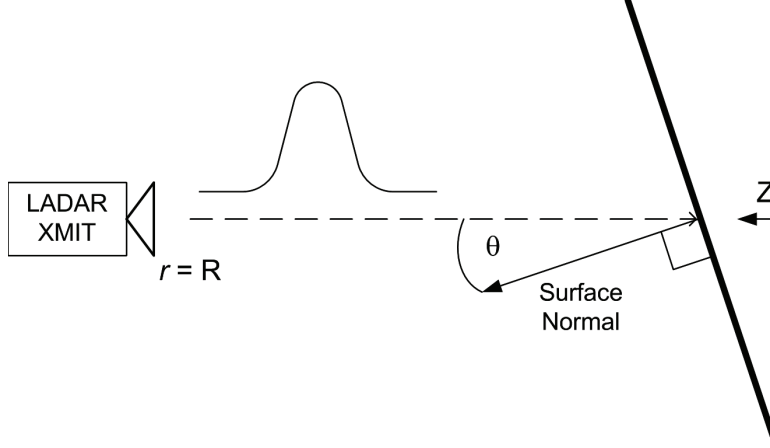


Figure 4.18. System diagram, showing target aspect angle with respect to laser boresight.




Figure 4.19. Sample 3rd phase image, showing fit-plane inclination angle. Angles are measured in degrees from the z -axis.

vector. The process is continued for the set of regions \mathcal{R}_n in each of the n range density bins. What results is a complete third-phase image of the scene, consisting of planar inclination angle, much like the image in Fig. 4.19. In this visualization, features which were once distinct, such as in the range image of Fig. 4.6, are markedly less so. For example, the foreground target and the fence in the background are nearly parallel.

Using the Canny edge detection filter on Fig. 4.19 yields the edge map in Fig. 4.20(a), where it can be seen that the near-parallel objects are too similar in value to be discerned at the 0.125 sensitivity threshold value, the default value. As the threshold is reduced to 0.01, in Fig. 4.20(b), the near-parallel object edges appear, but at the expense of identifying noisy, background objects edges as well. Since the Canny method is pre-disposed to finding weak edges, the fact of having to reduce to the threshold even further to select the edges in question indicates how subtle the transition between the regions truly is.




Figure 4.20. Edges of Fig. 4.19 using Canny edge filter.

Other published methods of finding surface normals (and similar local statistics) use sliding windows across local neighborhoods of pixels [1] [60]. With respect to the present discussion, sliding windows are not an ideal solution for calculating surface normals. The most significant drawback of the sliding window is the “leakage” of the window into other regions; this causes edge distortion in the third-phase image, especially when the relative difference of values in two adjacent regions is significantly large. By using KDE-based range segmentation in the initial steps, the edge corruption is avoided. The next section continues with the generation of third-phase images, describing the method for estimating material reflectance α .

4.3.2 Material Reflectance.

In this subsection, we discuss how the material reflectance of each object is estimated from the dataset. The transmitter and receiver of the ASC Portable Ladar Camera are co-located, so the system is monostatic. This means that the reflected light focused at the ladar detector is the backscattered reflectance from the illuminated targets. Hebert has previously modeled the backscattered laser intensity using a diffuse Lambertian model [36], as have Jutzi and Stilla [48]. Accounting for reduced intensity due to backscatter is an important contribution to modeling pulses reflected from resolved (multi-pixel) targets. The ladar beam is assumed to be sufficiently uniform across the selected region of the image; this choice is based on the ASC Portable Ladar Camera manufacturer’s specifications that the beam is uniform across the 3° FOV. Additional characterizations of the ladar beam and beam spreading effects are beyond the scope of this research.

The effective gain G_{eff} (or pulse amplitude) of the peak pulse in each waveform is functionally dependent on gain of the system \mathfrak{G} , as well as the surface reflectance α and incidence angle θ_i :

$$G_{eff} \propto \mathfrak{G} \alpha \cos \theta_i \quad (4.5)$$

where $0 \leq \alpha \leq 1$. Without an absolute, ground truth reference, we cannot determine the true value of α for all wavelengths. However, within each image, it is possible to estimate relative values of α using the gain and angle estimates that have previously been derived. The gain estimate, such as the one derived in (5.44) is thresholded at each pixel to remove any spurious (or negative) values; these values, should they exist, are replaced with the median of the ML gain estimate data. The “system gain” estimate, $\hat{\mathfrak{G}}$, is then calculated at every point as

$$\hat{\mathfrak{G}} \alpha \propto \frac{\hat{G}}{\cos \hat{\theta}_i^k} \quad (4.6)$$

where $\hat{\theta}_i^k$ is the estimated incidence angle for each segmented region r_k and \hat{G} is the gain estimator. A histogram of $\hat{\mathfrak{G}}$ is shown in Fig. 4.21(a), with the kernel density estimate shown in Fig. 4.21(c). Since the data is so heavily skewed, we calculate the 10% trimmed mean, and replace all the data beyond 2 standard deviations with the trimmed mean value. Trimmed mean filters have previously been used where images are corrupted with impulse noise [59]. In this instance, the outlier data represents impossibly high estimates of the gain estimate, which result from angle estimates that approach 90° .

The trimmed histogram is shown in Fig. 4.21(b), with the trimmed mean line indicated, as is the first and second standard deviation from the mean; the density estimate of the trimmed data is shown in Fig. 4.21(d). The third-phase image of the reflectance estimate is shown in Fig. 4.22. Similarly reflective materials (such as the plywood targets) should appear to be the same color. From a visual inspection, we see the disparate estimates of the material reflectance of both plywood targets. However, the front board and the background fence have similar reflectance. The background (free space) and foreground (grass) indicate almost negligible reflectance, since these objects are illuminated at a grazing angle.

In the following section, the example datasets from Section 4.1 are used to generate the




Figure 4.21. System gain estimate $\hat{\mathfrak{G}}$ density graphs. (a) Untrimmed histogram (b) 10% trimmed mean histogram (c) Untrimmed density and (d) 10% trimmed mean density.




Figure 4.22. Relative α estimate.

surface normal and material reflectance third-phase images.

4.3.3 Examples.

The example data from Sec. 4.1.4 is presented again to illustrate the generation of third-phase images. All the material reflectance images presented are derived using the

Table 4.4. Calculated Inclination Angles.


| | Target 1 | Target 2 | Fence |
|---------|----------|----------|-------|
| Trial 1 | 13° | 10° | 13° |
| Trial 2 | 8° | 29° | 20° |
| Trial 3 | 8° | 34° | 20° |
| Trial 4 | 8° | 60° | 20° |

10% trimmed mean of the \hat{G}_{MLE} data. The calculated inclination angles of the prominent targets in all the trials are shown in Tab. 4.4. In the Trial 1 dataset, the three prominent objects were roughly parallel. In Fig. 4.23(a) the targets are almost indistinguishable, given that the inclination angles of the targets are 13°, 10°, 13° respectively. The reflectance image in Fig. 4.23(b) shows, as would be expected, that the targets have very similar reflectance characteristics.

With Trial 2, the second board target is intentionally tilted, and the physical location of the ladar camera was slightly adjusted. For this reason, the angles of the fixed targets in the first 2 trials are different. The inclination angle of the boards in Fig. 4.24(a) are 8°, 29°, 20° respectively. As a result, the reflectance estimate in Fig. 4.24(b) shows that the estimate of the reflectance of the tilted board is less than the fixed boards.

In Trial 3, the second board is tilted again to 34°. The angle estimate of Fig. 4.25(a) influences the estimation of the material reflectance in Fig. 4.25(b), and the result is similar to that of Trial 2.

In the final trial, Trial 4, the second board is at the maximum tilt of 60°, which is seen in the angle image of Fig. 4.26(a). Observing the material reflectance estimate in Fig. 4.26(b) indicates that the estimates are not as poor as one would expect.

**Figure 4.23. Third-phase Images of Trial 1 data**




Figure 4.24. Third-phase Images of Trial 2 data.




Figure 4.25. Third-phase Images of Trial 3 data.




Figure 4.26. Third-phase Images of Trial 4 data.

4.3.4 Summary.

Third-phase ladar images are generated using data from multiple pixel waveforms. Certain bits of information, such as the object surface normal vector, or the material reflectance parameter, can only be found by analyzing the first- and second-phase images and transforming the data into a new visualization. The distinctive third-phase image is found to be extremely helpful for analyzing the characteristics of the underlying targets illuminated by the ladar camera, as has been shown throughout the section.

4.4 Chapter Summary

In this chapter, the proposed ladar segmentation method was introduced and demonstrated. A short discussion of the nature of laser radar images was presented, and this motivated the necessity of classifying ladar images into distinct “phases” of imagery. The segmentation method was described in Sec. 4.1. The method is initialized using non-parametric probability density estimation, as this classifies the data into probability range bins, which is the first step in the segmentation. The method is shown to have similar results to an iterative region-growing segmentation technique, which was adapted from contemporary research. Plane fitting of each segmented region is accomplished using principal components analysis, from which the plane inclination angle with respect to the ladar boresight is estimated. Also, the new third-phase image containing the relative material reflectance of the segmented object is estimated. A discussion of the strengths and weaknesses of the segmentation method was also presented.

V. Backscatter Reflectance Waveform Model

This chapter establishes and justifies the improved ladar signal model used throughout the remainder of the research. Drawing on hitherto accepted ladar signal models proposed in [14], [45], the model developed in this chapter unites the previous work with the simplest form of a bidirectional reflectance distribution reflectance (BRDF) model. Including a notional BRDF as part of a ladar system study is not a unique proposition [48], [92]. However, the choice of pulse shape, noise model, and evaluation methods in this and later chapters constitute a significant contribution to the overall body of knowledge.

The validity of the Poisson-distributed signal model has been discussed in Chapter III. While several temporal pulse models have been proposed by various other sources, the Gaussian pulse model is adopted for this research due to its well-known properties. This and other pulse models are parameterized by choices of gain (amplitude), target range (time to target), signal bias (additive noise), and pulse width (or length). These parameters may be estimated either independently or jointly, and the confidence in the estimates may be quantified by calculating a lower bound on the mean-square error (MSE) of the parameter estimate.

In Sec. 5.1 the Gaussian waveform model of Chapter III is modified to include the effects of diffuse backscatter. In Sec. 5.2 the parameters are evaluated using the Cramér-Rao Bound. A discussion of the effectiveness of alternative error bounds is provided in Sec. 5.3. The improvement in the lower error bound is quantified in this section. Finally, the reflectance fidelity pulse model is applied to experimental data in Sec. 5.4.

5.1 Backscatter Reflectance

Designing increasingly more accurate ladar waveform models has been the object of researchers for many years [70] [36] [48]. Accounting for the reflectance of objects can be approximated by integrating optical principles is the topic of this section. Investigation in this area was motivated by the analysis of Johnson [45], who modeled pulse expansion effects of non-resolved tilted targets. In order to model pulse expansion effects of resolved

targets, however, it is necessary to characterize the surface reflection using the bi-directional reflectance distribution function BRDF.

The bi-directional reflectance can be generally defined as the ratio of the radiance L scattered into the direction described by the orientation angles (θ_r, ϕ_r) to the irradiance E onto the reflecting surface from the (θ_i, ϕ_i) direction, i.e.,

$$f_{BRDF} = \frac{L(\theta_r, \phi_r)}{E(\theta_i, \phi_i)} \quad (5.1)$$

which has units of inverse steradians (sr^{-1}) [86]. The BRDF is also dependent on wavelength, but this dependence will be suppressed in the current analysis.

The BRDF has been used to model the radar cross section of various shaped solid objects illuminated by a radar [92]. An approximate form of the BRDF, based on the Torrance-Sparrow model, the sum of both reflected components (specular and diffuse), was proposed and verified in that paper. Since very few material-specific BRDF databases exist, researchers must make do with approximate forms of the function. For the present undertaking, since the transmitter and receiver are co-located, we may limit ourselves to being concerned about the “hot spot” backscatter reflection, no matter what type of object is being illuminated. The reflected radiance L_r at the backscatter angle can then be approximated in terms of the incident radiance L_i by

$$L_r \propto \alpha L_i \cos^2 \theta_i \quad (5.2)$$

where α is an experimentally derived constant and θ_i is the incident angle with respect to the surface normal [86]. Jutzi and Stilla [50] have suggested that the observed reflectance in the backscatter direction depends on the wavelength-specific reflectance ρ_m and the θ_i :

$$\rho_{diffuse} \propto \rho_m \cdot \cos(\theta_i) \quad (5.3)$$

The inclusion of the BRDF as part of the waveform amplitude function G means that the parameter estimates are also dependent on the reflectance α and the incidence angle

θ_i . The effective gain G_{eff} (or pulse amplitude) is functionally dependent on gain of the system \mathfrak{G} , as well as the surface reflectance α and incidence angle θ_i :

$$G_{eff} = \mathfrak{G} \alpha \cos \theta_i \quad (5.4)$$

The reflectance model proposed by Nitzan [70] is offered in comparison. The contours of effective gain for a fixed reflectance ($\alpha = 0.8$) are shown in Fig. 5.1. As the incidence angle approaches a grazing angle, the effective gain rapidly drops off. From this plot, the ambiguity between effective gain and tilt-induced gain is illustrated. To the observer, the effect of the BRDF term is manifested within the observed effective gain. However, by employing other tactics, such as plane fitting and surface normal calculation, the incidence angle of the target can be calculated, from which the true signal gain and material reflectance can be calculated.




Figure 5.1. Contours of effective gain $G_{eff} \propto (\mathfrak{G} \alpha \cos \theta_i) / (\sqrt{2\pi} \sigma_R)$. Reflectance parameter $\alpha = 0.8$.

Incorporating the BRDF into the standard pulse model of (3.8) requires a scaling of the signal power of the reflected waveform by the RHS of (5.2). Therefore, the simple temporal

pulse model can be written as

$$\begin{aligned}\lambda_{recv}(t) &= \frac{\mathfrak{G}\alpha \cos \theta_i}{\sqrt{2\pi}\sigma_R} \exp\left(\frac{-1}{2\sigma_R^2} \left(t - \frac{2R}{c}\right)^2\right) + B \\ &= G \exp\left(\frac{-1}{2\sigma_R^2} \left(t - \frac{2R}{c}\right)^2\right) + B\end{aligned}\tag{5.5}$$

where G is the effective amplitude gain of the Gaussian pulse shape. The pulse model in (5.5) will be referred to throughout as the backscatter reflectance waveform (BRW) model, as distinct from the standard model described by (3.8). Comparing the sample pulses from (3.8) and (5.5) in Fig. 5.2 reveals the amplitude difference between the two models. The reflectance term in this example is $\alpha = 0.7$, and the incidence angle θ_i is 25° .




Figure 5.2. Sample Gaussian pulse models from Eq. (3.8) (Standard) and Eq. (5.5) (BRW). Inclination angle $\theta_i = 25^\circ$, and $\alpha = 0.7$.

To demonstrate the backscatter reflectance behavior, the peak pulse amplitude of the standard model (3.8) is compared to the peak pulse amplitude of the backscatter reflectance model (5.5). The amplitude is calculated for incidence angles beginning at $\theta_i = 0^\circ$, incremented in steps of 15° up to $\theta_i = 75^\circ$. These comparative plots are shown in Fig. 5.3, with the simulated difference between the models plotted as the dotted line.




Figure 5.3. Normalized peak pulse amplitude of standard and BRW pulse models. Pulse spreading effects are ignored.


A comparison of the pulse power between the BRW and standard model is also instructive. As has already been discussed, the backscatter energy from normal incidence (which is the specular reflection) does not correspond to a $\alpha \cos \theta_i$ scaling. The pulse amplitude scaling introduced by (5.1) reduces the signal energy of the pulse reflecting from a non-zero incidence angle.

The power in each pulse can be found by integrating the received digital signal over time. Since the received lidar pulse is given by $\lambda(t|\mathbf{A}) = S(t, \mathbf{A}) + B$, where the signal part $S(t, \mathbf{A})$ is a Gaussian pulse, then the energy in each pulse with


$$\begin{aligned}
 E_{tilt} &= \int_{-\infty}^{+\infty} S(t, \mathbf{A}) dt \\
 &= \int_{-\infty}^{+\infty} \frac{\mathfrak{G} \alpha \cos \theta_i}{\sqrt{2\pi\sigma^2}} \exp \left\{ \frac{-(t - t_R)^2}{2\sigma^2} \right\} dt \\
 &= \mathfrak{G} \alpha \cos \theta_i
 \end{aligned} \tag{5.6}$$

The pulse energy in the standard model is simply $E_{std} = \mathfrak{G}$. In Fig. 5.4, the energy of a 1 ns and 3 ns pulse is calculated using both waveform models. The conversion from observed

photos-counts to power is assumed to be trivial. The pulse energy plots are an advantage over the amplitude plots of Fig. 5.3, since it can be seen that the reflected pulse energy from the standard model is assumed to be constant.



(a) 1 ns transmitted pulse width



(b) 3 ns transmitted pulse width

Figure 5.4. Signal pulse energy with respect to incidence angle. Comparison is made between BRW and Standard models of pulse amplitude.

5.1.1 Waveform Parameter Estimates.

Fitting the observed waveform to a model can be accomplished in several ways. In this section, waveform parameter estimates are generated using maximum-likelihood (ML) estimation MLE and correlation (or matched filtering). The correlation methods are shown to give similar results for both parametric and non-parametric formulations. All parameter estimates are discussed in relation to each other, and the estimator error is evaluated using the CRB.

The method of maximum likelihood is a commonly used technique for deriving parameter estimates in a signal. Let \mathbf{X} be a sample of data point, parameterized by $\boldsymbol{\theta}$, with joint likelihood function $L_{\mathbf{X}}(\boldsymbol{\theta})$. Then $\widehat{\boldsymbol{\theta}}_{ML}$ is the parameter value at which $L_{\mathbf{X}}(\boldsymbol{\theta})$ attains its global maximum, as a function of $\boldsymbol{\theta}$. Since, at a maximum, the gradient of the log-likelihood function vanishes, it is often advantageous to solve

$$\frac{dL_{\mathbf{X}}(\boldsymbol{\theta})}{d\boldsymbol{\theta}} = \mathbf{0} \quad (5.7)$$

instead, depending on the functional form of $L_{\mathbf{X}}(\boldsymbol{\theta})$. Note that the solutions to (5.7) are only *candidates* for the MLE; the global maximum condition must still be satisfied before the MLE is declared. In the following sections, estimates for range, gain, and bias will be derived. The range estimate is found using matched filtering, while the gain and bias are derived from an ML solution.

5.1.1.1 Range Estimator.

The parametric cross-correlation algorithm of [13] computes the Pearson-product correlation coefficient for different pairs of range and pulse width, and the combination that yields the highest overall correlation value ρ is selected for the joint estimate:

$$\hat{\rho}(\hat{R}, \hat{\sigma}) = \sum_{k=1}^K \frac{(X_k - \mu_X)(Y_k - \mu_Y)}{\sigma_X \sigma_Y} \quad (5.8)$$

where \hat{R} and $\hat{\sigma}$ are the final range and pulse width estimates, respectively, X_k is the k^{th} measured waveform sample, k being an integer, Y_k is the waveform model, σ_X is the square root of the total power in the waveform samples, σ_Y is the square root of the total power in the waveform model samples, and K is the total number of samples in both the measured and modeled waveform. The sample mean of X and Y is denoted by μ . If the pulse width is assumed to be constant, then the dimensionality of (5.8) is reduced to unity.

5.1.1.2 Gain Estimator.

An estimator for the waveform gain parameter can be obtained using the sum of the data observations. The derivation of this estimator is provided below, and is adapted from a similar estimator developed by Peterson [75].

Let a sequence of K observations of the time-varying Poisson process (3.7) be given by the data vector $\mathbf{Y} = [y_1, y_2, \dots, y_K]$. Each y_i is an observation of a Poisson random variable $Y_i \sim POI(\lambda(t_i))$. The time-varying mean of Y_i corresponds to the waveform components given in (3.8), where the exponential term describes the pulse shape, G is the effective pulse amplitude and B is the signal bias. Let the test statistic $T(\mathbf{y})$ be defined as the sum of the observations y_i :

$$T(\mathbf{y}) = \sum_{i=1}^K y_i \quad (5.9)$$

This statistic is equivalent to an observation from the random variable Z , defined as the sum of K Poisson random variables

$$Z = Y_1 + Y_2 + \dots + Y_K \quad (5.10)$$

Since each Y_i is distributed Poisson with parameter $\lambda(t_i)$, the random variable Z also has a Poisson distribution, with parameter given by

$$\begin{aligned} \lambda_Z &= \lambda(t_1) + \lambda(t_2) + \dots + \lambda(t_K) \\ &= \sum_{i=1}^K \lambda(t_i) \end{aligned} \quad (5.11)$$

This is a well-known result. The expectation of Z is therefore

$$\begin{aligned}
E(Z) &= \sum_{i=1}^K \lambda(t_i) \\
&= \sum_{i=1}^K (S(t_i) + B) \\
&= KB + G \sum_{i=1}^K \exp\left(\frac{-1}{2\sigma_R^2} \left(t_i - \frac{2R}{c}\right)^2\right)
\end{aligned} \tag{5.12}$$

Let the sum of the exponential term be given by

$$s_K = \sum_{i=1}^K \exp\left(\frac{-1}{2\sigma_R^2} \left(t_i - \frac{2R}{c}\right)^2\right) \tag{5.13}$$

which allows (5.12) to be shorted to

$$E(Z) = KB + Gs_K \tag{5.14}$$

Using this result, the functional form of the pdf of Z is

$$P(Z = z) = \frac{[KB + Gs_K]^z}{z!} \exp\{- (KB + Gs_K)\} \tag{5.15}$$

where $z = T(\mathbf{y}) = \sum_{i=1}^K y_i$. Taking the natural log of (5.15) gives the function

$$L_Z(z) = z \log(KB + Gs_K) - (KB + Gs_K) - z! \tag{5.16}$$

This is the log-likelihood for a single observation of z , i.e. the sum of the data of a single waveform. To obtain an estimate for G which maximizes (5.16), we first find the partial derivative of $L_Z(z)$ with respect to G :

$$\frac{\partial L_Z(z)}{\partial G} = \frac{zs_K}{KB + Gs_K} - s_K \tag{5.17}$$

Setting this equal to zero and solving for G gives the result

$$G = \frac{z - KB}{s_K} \quad (5.18)$$

Examining the second partial derivative,

$$\frac{\partial^2 L_Z(z)}{\partial G^2} = \frac{-zs_K^2}{(KB + Gs_K)^2} \quad (5.19)$$

if z, s_K, K, N, G are all strictly positive, then (5.19) is strictly negative. Then the gain estimate of (5.18) maximizes (5.16). The gain estimate is therefore

$$\hat{G} = \frac{z - KB}{s_K} \quad (5.20)$$

Taking $E(\hat{G})$, it can be shown that \hat{G} is an unbiased estimate for G . The next section details a method for finding a signal bias estimator.

5.1.1.3 Bias Estimator.

To find an estimate for the signal bias, the log-likelihood for Z , (5.16), is used again, but is maximized instead for B . The first partial derivative is

$$\frac{\partial L_Z(z)}{\partial B} = \frac{zK}{KB + Gs_K} - K \quad (5.21)$$

which, when set to zero, can be solved for B :

$$B = \frac{z - Gs_K}{K} \quad (5.22)$$

The second partial derivative with respect to B is strictly negative, so (5.22) is an estimator for bias which maximizes (5.16).

$$\hat{B} = \frac{z - Gs_K}{K} \quad (5.23)$$

Taking $E(\hat{B})$, it can be shown that \hat{B} is an unbiased estimate for B . Since calculation of (5.23) requires knowledge of the signal gain G , and this cannot be estimated simultaneously in practice, an alternate estimator can be derived using rank statistics of \mathbf{Y} . First, it is assumed that neither the first nor last element of \mathbf{Y} contain energy from the pulse shape. A global estimate for signal bias is found using the mean of y_1 and y_K at every pixel in the cuboid. If the cuboid has (x, y, z) dimensions of $(x_{max} \times y_{max} \times K)$, then the estimate for bias is

$$\hat{B} = Median(\mathbf{B}) \quad (5.24)$$

where

$$\mathbf{B} = \left\{ \frac{1}{2} (y_1(x, y) + y_K(x, y)) \right\}, \quad \begin{array}{l} x \in [1, x_{max}] \\ y \in [1, y_{max}] \end{array} \quad (5.25)$$

The estimate for \hat{B} is global median of the average of the first and last observations in each pixel of the cuboid. The methodology for obtaining the estimator is reasonable, since the data was collected with the express intent of avoiding pulse echoes at the edges of the range gate. Using the mean estimate would give greater influence to outlier values of \mathbf{B} , which is why the median operator was chosen instead.

5.2 Minimum Error Bounds

This section presents the derivation of the Cramér-Rao least error bound. For a more in-depth treatment, the reader is directed to the Appendix A, or to [98]. Given a single pulse shape parameter A , the variance on the error of an unbiased estimate \hat{a} of that parameter is bounded below by

$$VAR(\hat{a} - A) \geq \left\{ -E \left(\frac{\partial^2 L_{\mathbf{Y}}(A)}{\partial A^2} \right)^2 \right\}^{-1} \quad (5.26)$$

where the first and second derivatives of the likelihood function,

$$\frac{\partial L_{\mathbf{Y}}(A)}{\partial A}, \quad \frac{\partial^2 L_{\mathbf{Y}}(A)}{\partial A^2} \quad (5.27)$$

exist and are absolutely integrable. This inequality is commonly referred to as the Cramér-Rao bound (CRB) for an unbiased estimator [98].

Generalizing to simultaneous estimates of a vector of M parameters is straightforward. The unbiased estimators $\hat{\mathbf{A}} = [\hat{a}_1, \dots, \hat{a}_i, \dots, \hat{a}_M]^T$ are bounded below by

$$VAR(\hat{a}_m - A_m) \geq (\mathbf{J})_{mm}^{-1} \quad (5.28)$$

where $(\mathbf{J})_{mm}^{-1}$ is the mm th element in the inverse of the $M \times M$ Fisher Information Matrix (FIM) \mathbf{J} . The elements of \mathbf{J} are:

$$\begin{aligned} J_{ij} &= E \left(\frac{\partial L_{\mathbf{Y}}(A)}{\partial A_i} \frac{\partial L_{\mathbf{Y}}(A)}{\partial A_j} \right) \\ &= -E \left(\frac{\partial^2 L_{\mathbf{Y}}(A)}{\partial A_i \partial A_j} \right) \end{aligned} \quad (5.29)$$

which are the second partial derivatives of (5.32).

5.2.1 Likelihood Based on Product.

The function $L_{\mathbf{Y}}(\mathbf{A})$ is the joint log-likelihood function of the collected data at each pixel (i, j) in the detector array. To derive it, recall that the photocount at each pixel in the array can be modeled as a Poisson random variable $Y(t)$, with time-varying mean $\lambda(t|\mathbf{A})$, where \mathbf{A} is the vector of waveform parameters. Referring to (3.5), the joint probability of observing k photocounts at each time sample t_i is

$$P(Y_i = k|\mathbf{A}) = \frac{1}{n!} \left[\int_{t_i}^{t_i+\Delta t} \lambda(t|\mathbf{A}) dt \right]^n \exp \left\{ - \int_{t_i}^{t_i+\Delta t} \lambda(t|\mathbf{A}) dt \right\} \quad (5.30)$$

The joint probability (or likelihood function) for of a sequence of K events is then

$$\begin{aligned}
P[Y_1 = y_1, \dots, Y_K = y_K] &= p[y_1, \dots, y_K] \\
&= \prod_{i=1}^K \frac{\lambda(t_i|\mathbf{A})^{y_i}}{y_i!} \exp\{-\lambda(t_i|\mathbf{A})\} \\
&= \exp\left\{-\sum_{i=1}^K \lambda(t_i|\mathbf{A})\right\} \prod_{i=1}^K \frac{\lambda(t_i|\mathbf{A})^{y_i}}{y_i!}
\end{aligned} \tag{5.31}$$

and taking the natural log of (5.31), it follows that the log-likelihood function $L_{\mathbf{Y}}(\mathbf{A})$ is

$$L_{\mathbf{Y}}(\mathbf{A}) = \sum_{i=1}^K y_i \log \lambda(t_i|\mathbf{A}) - \sum_{i=1}^K \lambda(t_i|\mathbf{A}) - \sum_{i=1}^K \log(y_i!) \tag{5.32}$$

The equation in (5.32) calculates the joint likelihood of the observations in a single waveform. Parameterizing (5.32) using $\mathbf{A} = [R, G, B]$, and solving for the elements of the 3×3 FIM, as described by (5.29), (see details in Appendix A), the results for the diagonal elements of \mathbf{J} are:

$$\begin{aligned}
J_{11} &= \frac{4}{c^2} \frac{1}{\sigma_R^4} \sum_{k=1}^K \left(t_k - \frac{2R}{c}\right)^2 \left(\lambda(t|\mathbf{A}) - 2B + \frac{B^2}{\lambda(t|\mathbf{A})}\right) \\
J_{22} &= \frac{1}{G^2} \left[T_{\lambda}(A) - K2B + B^2 \sum_{k=1}^K \frac{1}{\lambda(t|A)} \right] \\
J_{33} &= \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})}
\end{aligned} \tag{5.33}$$

The off-diagonal elements of \mathbf{J} are:

$$\begin{aligned}
J_{12} = J_{21} &= \frac{2}{c} \frac{1}{\sigma_R^2} \frac{1}{G} \sum_{k=1}^K \left(t_k - \frac{2R}{c}\right) \left(\lambda(t|A) - 2B + \frac{B^2}{\lambda(t|A)}\right) \\
J_{13} = J_{31} &= \frac{2}{c} \frac{1}{\sigma_R^2} \sum_{k=1}^K \left(t_k - \frac{2R}{c}\right) \left(1 - \frac{B}{\lambda(t|A)}\right) \\
J_{23} = J_{32} &= \frac{1}{G} \sum_{k=1}^K \left(1 - \frac{B}{\lambda(t|A)}\right)
\end{aligned} \tag{5.34}$$

It is not necessary to find a closed form solution to these equations, because they can all be evaluated numerically. In Fig. 5.5, the time-amplitude plots of the six unique elements

Table 5.1. Waveform parameters for the sample plots of Fig. 5.5.

| Parameter | Value |
|----------------|-------|
| α | 1 |
| θ_i | 0 |
| σ_R | 10ns |
| \mathfrak{E} | 50 nJ |
| B | 1 |
| R | 208m |

of \mathbf{J} are shown. The parameters used to create these plots are given in Tab. 5.1. Referring again to Fig. 5.5, the diagonal elements ($J_{11}, J_{22}, \& J_{33}$) have even symmetry about the time-to-target, as do the gain-bias derivatives (J_{23}, J_{32}). The two plots with odd symmetry, corresponding to the range-gain derivatives (J_{12}, J_{21}) and range-bias derivatives (J_{13}, J_{31}), have odd symmetry around the time-to-target. The sum, therefore, of these functions given by (5.34), is approximately zero.

Given that four of the elements of \mathbf{J} vanish, the FIM can be written as

$$\mathbf{J} = \begin{bmatrix} J_{11} & 0 & 0 \\ 0 & J_{22} & J_{23} \\ 0 & J_{32} & J_{33} \end{bmatrix} \quad (5.35)$$

Bounds on estimates are given by the elements of the inverse \mathbf{J} matrix. To find the inverse, observe that \mathbf{J} is a block-diagonal matrix of the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{B} & 0 \\ 0 & \mathbf{C} \end{bmatrix} \quad (5.36)$$

The inverse of \mathbf{A} is given by

$$\mathbf{A}^{-1} = \begin{bmatrix} \mathbf{B}^{-1} & 0 \\ 0 & \mathbf{C}^{-1} \end{bmatrix} \quad (5.37)$$

Matching the matrix elements of (5.36) to the elements in (5.35), the corresponding inverse

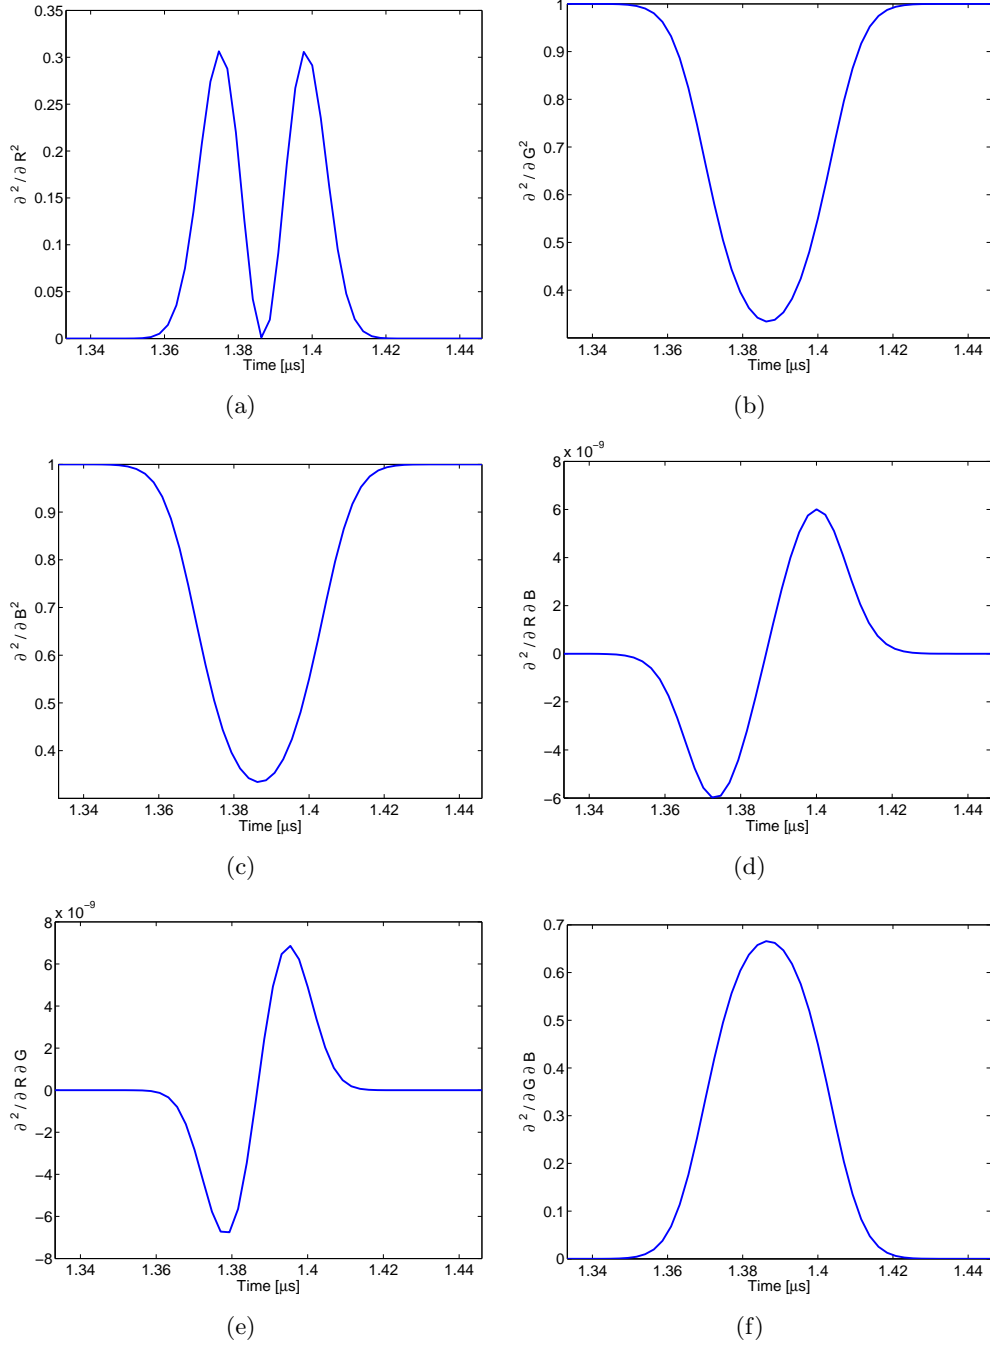


Figure 5.5. Time-amplitude plots of the elements of J . (a) J_{11} (b) J_{22} (c) J_{33} (d) J_{13} (e) J_{12} (f) J_{23} . Parameters for the plots are shown in Tab. 5.1.

of J is

$$J^{-1} = \begin{bmatrix} J_{11}^{-1} & 0 & 0 \\ 0 & J_{33}/\det C & -J_{32}/\det C \\ 0 & -J_{23}/\det C & J_{22}/\det C \end{bmatrix} \quad (5.38)$$

where $\det \mathbf{C} = J_{33}J_{22} - J_{32}J_{23}$. Using (5.28) and (5.38), the least error bound for the joint estimate of $\mathbf{A} = [R, G, B]$ is obtained.

5.2.2 Likelihood Based on Sum.

The gain and bias estimators were derived using the likelihood of the sum of the observations $T(\mathbf{y}) = \sum_{i=1}^K y_i$. Neither range nor pulse shape information cannot be obtained from this statistic, so no bounds will be derived for estimators of those parameters. The log-likelihood function of the sum of the data is given by (5.16):

$$L_Z(z) = z \log(KB + Gs_K) - (KB + Gs_K) - z! \quad (5.39)$$

The unknown parameters in (5.39) are the gain G , and the bias B . If both parameters are estimated simultaneously, then the FIM is 2×2 with elements comprised of:

$$\begin{aligned} J_{11} &= \frac{s_K^2}{KB + Gs_K} \\ J_{22} &= \frac{K^2}{KB + Gs_K} \\ J_{12} = J_{21} &= \frac{s_K K}{KB + Gs_K} \end{aligned} \quad (5.40)$$

Since the denominator is common to all the entries, it is pulled outside, so \mathbf{J} is written as

$$\mathbf{J} = \frac{1}{KB + Gs_K} \begin{bmatrix} s_K^2 & s_K K \\ s_K K & K^2 \end{bmatrix} \quad (5.41)$$

In this form it is clearly seen that the second column is a scalar multiple of the first column, where the scaling factor is K/s_K ; this causes the determinant of the matrix vanishes. The matrix, therefore, is singular and the CRB for a simultaneous parameter estimate does not exist. To overcome this, the bound on each parameter will be derived separately, assuming clairvoyant knowledge of the other parameter.

Let (5.39) be parameterized by gain G . Then, the variance of an unbiased estimator for

G is bounded below by

$$VAR(\hat{G}) \geq -E\left(\frac{\partial^2 L_Z(G)}{\partial G^2}\right) = \frac{s_K^2}{KB + Gs_K} \quad (5.42)$$

Likewise, if (5.16) is parameterized by bias B , then the variance of an unbiased estimator for B is bounded below by

$$VAR(\hat{B}) \geq -E\left(\frac{\partial^2 L_Z(G)}{\partial G^2}\right) = \frac{s_K K}{KB + Gs_K} \quad (5.43)$$

Since information is lost during the summation process, the CRB derived using $T(\mathbf{y})$ will only be as good or worse than the CRB derived using (5.31); it cannot be superior. In the following sections, the relative performance of each CRB derivation will be examined.

5.2.3 Estimator Performance.

According to the ladar signal model of (3.8), there are four basic parameters that can be estimated from the detected waveform: target range (R), effective gain (G_{eff}), pulse length (σ_R), and signal bias (B). It has also been shown that G_{eff} is functionally dependent on the material reflectance parameter α and incidence angle θ_i , although estimators for these parameters were derived in Chapter IV. Throughout this section, it is assumed that the received waveforms are single-echo pulses, *i.e.* reflections from a single surface.


5.2.3.1 Range Estimate.

The ML estimate for pulse range is not calculated in closed form, since it is dependent on a numerical maximization. Instead, the waveform correlation operation described in (3.12) is used to obtain estimates of the target range. The simulation parameters are given in Table 5.2. The matched filter range grid spacing was set to $\delta R = 0.2$ m, so the true target range R was an exact multiple of δR . The signal was subjected to Poisson noise in a Monte Carlo simulation of 501 iterations. Both parametric (*i.e.* Pearson) and non-parametric


Table 5.2. Waveform parameters for the sample plots of Fig. 5.6.

| Parameter | Value |
|------------|-----------|
| α | 1 |
| θ_i | 0 |
| σ_R | 3 ns |
| B | 100, 1000 |
| R | 208m |
| δR | 0.2 m |

(*i.e.* rank) correlation were used to obtain range estimates, with the resulting RMSE/CRB comparison shown in Fig. 5.6. The two plots show the difference in range estimate due to variable signal noise B .



(a) Lower Bound on $R(B = 100)$ estimator



(b) Lower Bound on $R_n(B = 1000)$ estimator

Figure 5.6. CRB of range matched-filter estimate with (a) $B = 100$, and (b) $B = 1000$.

The parametric correlation estimate has, overall, a superior performance to the rank correlation when compared to the CRB. The bound suggests improved performance as the effective gain increases, but above a certain gain, the estimator performance ceases to improve and levels off. As pointed out in [83], this feature is indicative of the deterministic modeling error that can be found in many correlation-based estimator approaches.

5.2.3.2 Gain Estimate.

As the CRB is used to evaluate unbiased estimates, it is necessary to confirm this property for \hat{G} . The estimate for G is given in (5.20). Taking the expectation of \hat{G} ,

$$E(\hat{G}) = E\left(\frac{z - KB}{s_K}\right) \quad (5.44)$$

The expectation operates only the sum of the observations $z = \sum_{i=1}^K y_k$. Since z is an instance of the random variable Z , $E(z)$ is given by (5.14), so (5.44) reduces to

$$\begin{aligned} E(\hat{G}) &= \frac{E(z) - KB}{s_K} \\ &= \frac{KB + Gs_K - KB}{s_K} = G \end{aligned} \quad (5.45)$$

So the gain estimator is unbiased, and the CRB can be used to evaluate it. In Fig. 5.7, the estimator of the effective gain, parametric in signal bias, is compared to the CRB. Simulation parameters, including the Monte-Carlo iterations, are shown in Table 5.2. The gain estimate assumes knowledge of the bias, which is found using the bias estimator of (5.24). The gain estimate achieves the CRB of (5.42), but not the CRB of (5.33) (dotted line). When the signal bias is 100, the separation between the two bounds has a maximum of $10^{0.09}$ at the low-gain extreme, with the separation decreasing as gain increases; this is shown in Fig. 5.7(a). When the bias is increased to 1000, the separation between the bounds has a maximum of $10^{1.4}$, which is shown in Fig. 5.7(b). The bound on a uniformly random guess at gain is shown at the top of the graph.

5.2.3.3 Bias Estimate.

The estimate of the signal bias using the median operator is given in (5.24). A closed form expression to check the estimator bias cannot be derived, although the bias can be calculated numerically. However, the expectation of the estimator from (5.23) can be evaluated in closed form. Since (5.23) has nearly the same form as \hat{G} , the expectation of \hat{B}




Figure 5.7. CRB of gain estimate with (a) $B = 100$, and (b) $B = 1000$.

is easily found to be

$$\begin{aligned}
 E(\hat{B}) &= \frac{E(z) - Gs_K}{K} \\
 &= \frac{KB + Gs_K - Gs_K}{K} = B
 \end{aligned} \tag{5.46}$$

showing that this ML estimate for B is itself unbiased. The plot of the estimator and bounds in Fig. 5.8 show the performance of this estimator as well as the median estimator of (5.24). In addition, the CRB using the product of the data and the sum of the data are both shown. In both Fig. 5.8(a) and Fig. 5.8(b), the ML estimate (5.23) achieves the CRB based on the sum. The variance of the median estimator is on the order of 10^{1-3} , which, although it does not achieve either CRB, is constant with respect to gain. The variance of the ML estimate for bias in Fig. 5.8(a) does begin to increase as gain increases, but this is to be expected. Along with the noise variance analysis performed in Section 3.3.1, the evaluation of the estimator variance using the CRB confirms that the median estimator from (5.24) is acceptable.

5.3 Other Least-Error Bounds

The CRB has been the most widely used minimal mean-square error bound within the signal processing community for many years. This is due in no small part to its relative ease of calculation [51] [98]. However, the CRB is used only in the high-SNR (or asymptotic)

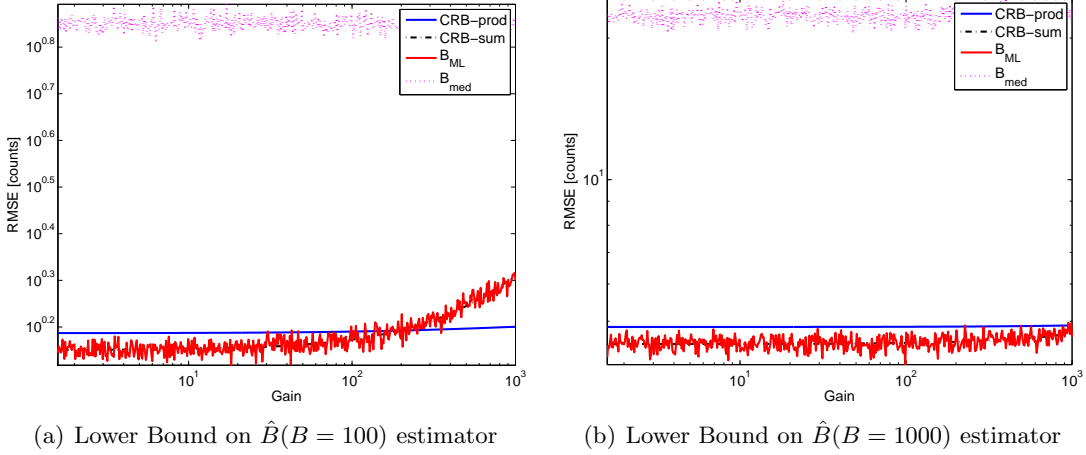


Figure 5.8. CRB of bias estimate \hat{B} with (a) $B = 100$, and (b) $B = 1000$.

region where the MSE is very small. As the SNR (or number of observations) decreases, information in the estimator is diminished, and its performance mimics that of a uniform random variable. This region is the “no information” area. Between the asymptotic and no information region is the threshold region. While the CRB is not appropriate for use in the threshold region, other minimal-MSE bounds have been developed that can be used. This section examines their use for the problem at hand.

5.3.1 Bobrovksy-Zakai Bound.

A tighter (*i.e.* closer to the MLE/MAP estimator) bound than the BCRB is the Bobrovksy-Zakai bound (BZB). A Bayesian bound, the BZB is a large-error bound on the MAP estimator. It is given by [81]

$$\zeta \geq \frac{h^2}{\iint \frac{p^2(\mathbf{x}, \theta + h)}{p(\mathbf{x}, \theta)} d\mathbf{x} d\theta - 1} \quad (5.47)$$

where the integral is taken over the entire observation space \mathcal{X} and parameter space Θ . The BZB is a maximization over the free parameter h , which makes it similar to a finite-difference derivative [100].

The integral in the denominator of (5.60) can be rewritten as:

$$\iint \frac{p^2(\mathbf{x}, \theta + h)}{p(\mathbf{x}, \theta)} d\mathbf{x} d\theta = \int_{\Theta} \frac{p^2(\theta + h)}{p(\theta)} \int_{\mathcal{X}} \frac{p^2(\mathbf{x}|\theta + h)}{p(\mathbf{x}|\theta)} d\mathbf{x} d\theta \quad (5.48)$$

Using the joint likelihood function from (5.31), we can write $\frac{p^2(\mathbf{x}|\theta + h)}{p(\mathbf{x}|\theta)}$ as

$$\frac{p^2(\mathbf{x}|\theta + h)}{p(\mathbf{x}|\theta)} = \prod_{i=1}^M \frac{Q_i(\theta + h)^{2x_i} \exp\{-2Q_i(\theta + h)\}}{Q_i(\theta)^{x_i} \exp\{-Q_i(\theta)\}} \frac{x_i!}{(x_i!)^2} \quad (5.49)$$

We make the following substitutions,

$$\begin{aligned} c_{1i} &= Q_i(\theta) \\ c_{2i} &= Q_i(\theta + h) \\ \phi_i &= \frac{c_{2i}^2}{c_{1i}} \end{aligned} \quad (5.50)$$

and then (5.49) becomes

$$\frac{p^2(\mathbf{x}|\theta + h)}{p(\mathbf{x}|\theta)} = \left[\prod_{i=1}^M \exp\{-2c_{2i} + c_{1i}\} \right] \left[\prod_{i=1}^M (\phi_i)^{x_i} \frac{1}{x_i!} \right] \quad (5.51)$$

where the first term has no dependence on \mathbf{x} . Then the integral over the observation space \mathcal{X} from (5.48) becomes

$$\int_{\mathcal{X}} \frac{p^2(\mathbf{x}|\theta + h)}{p(\mathbf{x}|\theta)} d\mathbf{x} = \left[\prod_{i=1}^M \exp\{-2c_{2i} + c_{1i}\} \right] \int_{\mathcal{X}} \left[\prod_{i=1}^M (\phi_i)^{x_i} \frac{1}{x_i!} \right] d\mathbf{x} \quad (5.52)$$

We recognize the kernel of the joint Poisson likelihood function in the integrand, which

results in

$$\begin{aligned}
&= \left[\prod_{i=1}^M \exp \{-2c_{2i} + c_{1i}\} \right] \int \prod_{i=1}^M \frac{(\phi_i)^{x_i}}{x_i!} \frac{\exp(-\phi_i)}{\exp(-\phi_i)} d\mathbf{x} \\
&= \left[\prod_{i=1}^M \exp \{-2c_{2i} + c_{1i} + \phi_i\} \right] \underbrace{\int \prod_{i=1}^M \frac{(\phi_i)^{x_i}}{x_i!} \exp(-\phi_i) d\mathbf{x}}_{=1} \\
&= \prod_{i=1}^M \exp \{-2c_{2i} + c_{1i} + \phi_i\} \\
&= \exp \left\{ \sum_{i=1}^M \phi_i - 2c_{2i} + c_{1i} \right\}
\end{aligned} \tag{5.53}$$

which is independent of \mathbf{x} . The Bobrovsky-Zakai bound is finally given by

$$\zeta \geq \frac{h^2}{\int_{\Theta} \frac{p^2(\theta+h)}{p(\theta)} \exp \left\{ \sum_{i=1}^M \phi_i - 2c_{2i} + c_{1i} \right\} d\theta - 1} \tag{5.54}$$

One important requirement of the BZB is that the prior distribution of the parameter cannot be defined over a finite interval [100]. If a beta prior is used, as was the choice in the previous section, then the BZB does not exist, and the bound collapses to the BCRB.

The Bayesian information within the integral makes this a very complicated expression to evaluate, depending on the choice of prior. For a simpler calculation, we can look at the Barankin bound (BRB), which is the deterministic equivalent of the Bobrovsky-Zakai bound.

5.3.2 Barankin Bound.

In essence, the Barankin bound can be found from the BZB by dropping the prior information from the expression (5.54), leaving

$$\xi \geq \frac{h^2}{\exp \left\{ \sum_{i=1}^M \phi_i - 2c_{2i} + c_{1i} \right\} - 1} \tag{5.55}$$

where the constants from (5.50) still apply. An analytic solution to the Barankin bound does not generally exist, so many researchers have proposed numerical methods. As with

the BZB, the BRB is optimized over a set of test points $\{h\}_{1 \times K}$. In general, the choice of the number of test points K is completely arbitrary [88]. A common technique is to simply take K points equally spaced in the parameter interval of interest, increasing K until the bound converges [64]; this is the approach taken here.

5.3.3 Bayesian Abel Bound.

We will show the Bayesian Abel bound on the range θ of the ladar pulse model. By combining the large and small error bounds, the tractable form of the Bayesian Abel bound (BAB) is tighter in the SNR threshold region than the Bayesian CRB (BCRB) and Bobrovsky-Zakai bound (BZB). It can be achieved with low computational cost when certain conditions are met [80]. When the BCRB and BZB are available, the single-order BAB is:

$$BAB_{1,1} = \sup_h \frac{BCRB^{-1} + BZB^{-1}(h) - 2\phi(h)}{BCRB^{-1}BZB^{-1}(h) - \phi^2(h)} \quad (5.56)$$

where $\phi(h)$ is given by

$$\phi(h) = \frac{1}{h} \int_{\Theta} \int_{\Omega} \frac{\partial \log p(\mathbf{y}, \theta)}{\partial \theta} p(\mathbf{y}, \theta + h) d\mathbf{y} d\theta \quad (5.57)$$

The derivation details of each bound are given in Appendix B. The Bayesian Cramér-Rao utilizes the *a priori* probability density of the parameter θ and provides a global bound that does not depend on the value of the parameter on a specific trial [99]. Assume the *a priori* probability density on θ to be normally distributed with mean μ_θ and variance σ_θ^2 :

$$\begin{aligned} p_\Theta(\theta) &\sim \mathcal{N}(\mu_\theta, \sigma_\theta^2) \\ &= \frac{1}{\sqrt{2\pi}\sigma_\theta} \exp\left(-\frac{1}{2\sigma_\theta^2}(\theta - \mu_\theta)^2\right) \end{aligned} \quad (5.58)$$

Alternatively, we could have chosen a uniform distribution on θ in the interval $[0, \theta_{max}]$, where θ_{max} is sufficiently large; a beta distribution would also be a useful choice, given its versatility. Consequently, using (5.58) the Bayesian Cramér-Rao bound (BCRB) on an

unbiased estimator of range θ is given by

$$\begin{aligned} \text{BCRB} &= \left(-E_{\mathbf{x},\theta} \left(\frac{\partial^2 L_D(\theta)}{\partial \theta^2} \right) \right)^{-1} \\ &= \frac{\sigma_\theta^2}{f_s \times \text{SNR} \frac{32\sigma_\theta^2}{3c^2\sigma_w} + 1} \end{aligned} \quad (5.59)$$

and the Bobrovsky-Zakai bound (BZB) is given by

$$\begin{aligned} \text{BZB} &= \sup_h \frac{h^2}{E_{\mathbf{x},\theta} \left[\left(\frac{p(\mathbf{x},\theta+h)}{p(\mathbf{x},\theta)} - 1 \right)^2 \right]} \\ &= \sup_h \frac{h^2}{e^{\frac{f_s \frac{64\text{SNR}h^2(12+c^2\sigma_w^2)}{3c^4\sigma_w^3} + \frac{h^2}{\sigma_\theta^2}} - 1}} \\ &= \sup_h \frac{h^2}{e^{h^2 \left(f_s \frac{64\text{SNR}(12+c^2\sigma_w^2)}{3c^4\sigma_w^3} + \frac{1}{\sigma_\theta^2} \right)} - 1}} \end{aligned} \quad (5.60)$$

where f_s is the sampling frequency of the ladar system. The Bobrovsky-Zakai bound must be optimized over the test point h , and converges to the BCRB in the limit as $h \rightarrow 0$ [100]. Note that the exponential term in the denominator of (5.60) can be simplified into the form $\exp(\alpha h^2)$; maximizing over h is a simple process. In general, however, it is common to choose the free parameter h , the test point, on the parameter support which is approximated by $[-3\sigma_\theta, 3\sigma_\theta]$ [81].

The single-order Bayesian Abel Bound ($\text{BAB}_{1,1}$) is therefore given by

$$\text{BAB}_{1,1} = \sup_h \frac{\text{BCRB}^{-1} + \text{BZB}^{-1}(h) - 2\phi(h)}{\text{BCRB}^{-1}\text{BZB}^{-1}(h) - \phi^2(h)} \quad (5.61)$$

where

$$\begin{aligned} \text{BCRB} &= \frac{\sigma_\theta^2}{f_s \times \text{SNR} \frac{32\sigma_\theta^2}{3c^2\sigma_w} + 1} \\ \text{BZB} &= \frac{h^2}{e^{h^2 \left(f_s \frac{64\text{SNR}(12+c^2\sigma_w^2)}{3c^4\sigma_w^3} + \frac{1}{\sigma_\theta^2} \right)} - 1} \end{aligned} \quad (5.62)$$

and

$$\phi(h) = \frac{1}{\sigma_\theta^2} + f_s 4 \left(\frac{8h^2}{3c^2\sigma_w} + \frac{\sigma_w}{5} \right) \times \text{SNR} \quad (5.63)$$

5.3.4 Simulation.

The parameters of a typical laser radar system were used in the signal model, found in Tab. 3.1. The time-varying Poisson rate was assumed to be a Gaussian pulse with variable gain parameter G_{eff} , a transmitted energy level of 1 nJ, and a pulse width of $\sigma_w = 5$ ns. The sampling rate of the detector was assumed to be 500 MHz, which results in a sampling period of $\Delta t = 2$ ns. The system captures $M = 21$ samples, which are spaced Δt ns apart, resulting in a time window t_d of 42 ns, and an equivalent range gate of about 12 m. It was assumed that the illuminated target was located exactly at the $fl(M/2) = 10$ th sample.

Because of the particular form of the BZB in (5.54), both it and the BAB (5.61) converge to the BCRB; this can be seen in Fig. 5.9. To maximize (5.54) over h , the denominator must be minimized, causing the exponential term $\exp\{\alpha h^2\}$ to approach unity. However, this process simultaneously forces $h \rightarrow 0$, which causes the overall bound to reduce to the Bayesian Cramér-Rao Bound.




Figure 5.9. The Bobrovksy-Zakai bound and Bayesian Abel Bound converge to the Bayesian CRB in the limit as the test point $h \rightarrow 0$.

5.3.5 Discussion.

In every case attempted, the tighter bound calculations collapsed to the Cramér-Rao bound. This is for two reasons: the pulse shape (linear, Gaussian) and the noise model. It was therefore decided to continue the pulse model analysis using the proven Cramér-Rao Bound analysis techniques. It is similar to a model put forth by Johnson in [45], but with the important addition of a bi-directional reflectance distribution function (BRDF) that more accurately models the amplitude of a pulse reflected from a tilted, diffuse surface. The BRDF is applied to pulse models in both Poisson- and Gaussian-noise simulations. Lower-bound expressions for estimates of the range, gain, and bias are re-created using the BRDF. Future attempts to re-derive the higher order bounds using assuming a negative binomial distribution of the photocount (vice Poisson) would be a valuable contribution.

5.4 BRWM Verification

The objective of the analysis is to link the measurement of the reflected waveform to the calculated incidence angle of each target board. It will be shown that the incidence angle is also related to the received pulse length and peak pulse amplitude. A summary of the analytical findings is presented at the conclusion of the section.

5.4.1 Incidence Angle Estimation.

The inclination angle calculation of the reference target indicates that while the first (fixed) board and fence are parallel, none of the boards were in fact oriented normal to the laser LOS. During data collection in the field, it was not practical to precisely align the boards with the ladar camera LOS at a distance of 150+ m. However, since only the second board was rotated between each measurement, the data are sufficient to use as reference points. In the next section, the amplitude variation across the planar surfaces as a function of the bias angle is discussed.




Figure 5.10. Ladar system diagram, showing target board tilted at inclination angle θ_i , with respect to ladar LOS.

5.4.2 Amplitude Variation.

As previously discussed, the backscatter signal collected by the ladar camera is of primary interest for this experiment. The reduction in signal amplitude as a function of both target reflectivity ρ and inclination angle θ_i has been studied for tightly controlled experiments, and simulations in other publications, such as [50].

The present analysis will verify the backscatter reflectance behavior predicted by (5.2) by examining the waveform characteristics of the pulses reflected from the targets described in the previous section. In each segmented region of the ladar image, the maximum pulse amplitude at each pixel is measured. Within each segmented region, the mean maximum pulse amplitude is then calculated, which represents the average maximum value for a randomly selected waveform reflected from one of the target surfaces. This process is repeated for each region in the image and for each target setup, yielding four data points for each board. The measured mean peak amplitude, with respect to incidence angle, for

the reference board is plotted against that of the rotated board in Fig. 5.11. Also shown for comparison is the $\alpha \cos \theta$ scaling of the reference board amplitude. The black dotted line is the measured difference between the reference board (red line) and the rotated board (blue line).

The first point of discussion is the peak pulse amplitude for the first (or reference) data point, corresponding to an inclination angle of $\sim 8^\circ$. At this point, the peak pulse amplitude for both the first (fixed) and second (rotated) board is nearly identical; that is, the amplitude does not appear to decrease at the $\alpha \cos \theta_i$ rate expected by (5.2). It is not until the second rotation angle, at $\sim 29^\circ$, that the amplitude of the rotated board matches the model prediction. The reason for this apparent disparity is that, in the first case, the inclination angle - and therefore backscatter angle - is nearly specular. For specular backscatter, the proposed BRW model breaks down, and waveform amplitude is affected only by surface reflectance α .




Figure 5.11. Measured mean peak amplitude of target boards with respect to laser incidence angle.

The reverse situation, estimating the incidence angle and/or reflectance parameter from the pulse amplitude, can only be accomplished if a reference waveform is available. Outside of a controlled environment, a reference waveform may not always be available. However,

since the experimental setup intentionally included a fixed reference object, this avenue is available for use. In the next section we discuss the pulse shape measurements.

5.5 Chapter Summary

In this chapter, an updated pulse model of a Poisson-distributed ladar signal was presented. The amplitude of a Gaussian-shaped pulse model was shown to be dependent on the incidence angle of the incoming beam, as well as the optical reflectance of the illuminated material. Parameter estimates for the range, gain, and signal bias were derived using a mixture of parametric and non-parametric matched filtering, as well as ML and other estimates. The RMSE of each parameter was evaluated using the CRB. The backscatter reflectance waveform (BRW) model signal was introduced consistent with other models of laser pulse backscatter. By making effective gain dependent on the BRDF, a more accurate pulse model is obtained than has been previously used. The estimate of target range and other waveform parameters was evaluated using the CRB. Finally, in Section 5.4, experimental data was evaluated to test the newly introduced waveform model.

VI. Ladar Feature Analysis

Robust object classification remains a top priority in many research fields including computer vision, target recognition, navigation, and remote sensing. The goal of distinguishing the objects in the FOV from the background is of value no matter which specific application is at hand. Ever-more technologically advanced sensors continually provide researchers with not only more, but higher fidelity data with which to accomplish their mission. One of the latest sensor devices to be investigated for usefulness is the 3-D flash laser radar camera. The 3-D flash ladar camera uses an array of detector pixels to obtain a complete frame (angle-angle-time) of ladar data from one laser pulse. Since each pixel can count time to the target, a full-waveform signal is recorded at each position in the detector array. In effect, the camera records a moving picture of the scene as the pulse travels across the range gate. This technology is distinct from scanning ladar systems in that the broad area, short laser pulse illuminates, or “flashes” the entire scene at once. One advantage of the flash technology is the ability to capture a target image in one frame, as opposed to the multiple rastered images obtained from a scanning laser system. By “flashing” the scene, the environmental effects (such as atmospheric turbulence, or platform motion) are more or less constant across the entire detector array.

In this chapter, nonparametric statistics are proposed to estimate the distributions of objects in 3-D ladar data images. Bootstrap resampling of the waveform parameters is used as part of the feature selection. We also discuss the use of bootstrapping the local statistics of the waveform parameters in certain cases. The outline of this chapter is as follows. Sect. 6.1 discusses the likely target and image features and how to obtain them. The feature selection is tested in Sec. 6.2 using both simulated and actual data. Finally, the conclusions are given in Sec. 6.3.

6.1 Feature Extraction

In Chapter V, the procedure for estimating ladar waveform parameters was examined in detail. Consisting of the range (\hat{R}), system gain or amplitude ($\widehat{G_{eff}}$), signal bias (\hat{B}),

and pulse length ($\widehat{\sigma_R}$), along with the total number of pulses (M), these parameters are the primary inputs to the proposed backscatter reflectance waveform model, and, in the case of the range estimate, are the basis of the image segmentation algorithm. A usable classification scheme, however, will not be limited to this data set, but will include features extracted using the segmentation as well. Again, as described in Chapter V, the third-phase data extracted from the segments include the object inclination angle $\hat{\theta}_i$ and the estimated system gain

6.1.1 Image Statistics.

Recent obstacle detection and avoidance algorithm development has emphasized the contribution of local point statistics in effective computer vision, especially in applications involving semi-autonomous motorized vehicles [60] [101]. In the present context of segmentation and classification of stationary ladar images, local statistics can also be implemented as part of a feature extraction algorithm. Local statistics are calculated in a local neighborhood of k pixels, where k is determined experimentally. The value of k (i.e. size of the neighborhood) can range from 1, which is the point statistic itself, up to the size of the segment in question. The optimum neighborhood in this methodology tends to be $k < 30$ pixels.

The statistics themselves are calculated using windows of size k . If the statistic is the local mean, then the window behaves as an averaging filter; a median filter would have very much the same implementation. For higher-order statistics, the physical meaning of the output is less obvious. The local range variance, for example, would be near zero for a surface oriented perpendicular to the ladar boresight, while a tilted surface would have a non-zero variance. Additional local statistics of skewness and kurtosis measure the weight of the distribution about the local sample mean [4] [60]. The window sizes for the higher order statistics usually tend to be larger than the first-order statistics, due to the nature of the measurements.

The objective of the classification algorithm is to measure features which capture the behavior of the illuminated objects in the image. Since the segmentation was initialized

using the image range data, it is appropriate to consider statistics of the intensity and reflectance as the classification features. Specifically, the local skewness of system gain (

6.1.2 Sample Density Estimation.

The method of kernel density estimation has already been used with great effect to accomplish range density segmentation in Chapter IV. Estimating the range probability density was useful for parsing the image into bins, since the total number of data points was very large. However, for classification feature analysis, the number of data points in each segment is typically much smaller than the number of data points used in the segmentation of the total range image. In addition, the calculation of higher-order statistical distributions of the data (*e.g.* kurtosis) tend to require a large sample size; the size of each segmented region may be insufficient for calculating these distributions.

One method that compensates for small sample size is bootstrap resampling. The bootstrap procedure draws a sample from the original data using sampling with replacement; the bootstrap sample is the same size as the original data, but is not identical. The desired statistic is calculated from the bootstrap sample, and the process is repeated several times. The shape of the resulting distribution of bootstrap statistics allows for inference to be drawn about the distribution of the desired statistic. Bootstrap resampling is also used to generate confidence intervals of the statistic in question, a scenario which commonly occurs when the size of the original sample is small.

If necessary, KDE can be used to smooth the distribution of the bootstrapped statistics. This would be helpful mainly for visualization reasons, since the statistic is generated with previously-segmented regions, and further parsing of the histogram is not generally required for feature selection.

6.2 Feature Selection Examples

This section illustrates the feature extraction method using both local statistics and bootstrap resampling. A simulation using a multi-faceted board in Sec. 6.2.1 demonstrates calculating the local statistics of segmented regions. In Sec. 6.2.2, data from field

experiments are used to generate bootstrap statistics; the content in this section is drawn from [12].

6.2.1 Multi-Facet Board Simulation.

In this section, a simulation of a multi-faceted board is used to demonstrate the feasibility of calculating local statistics for feature classification. The features are extracted from the previously introduced waveform parameters, and from the system gain.

A set of flat diffuse targets is the subject of the simulation. Four targets, each with material reflectance parameter $\alpha = 0.7$, are located at a minimum range of 156.7 m from the ladar camera. Material reflectance was simulated using a normal distribution, to account for variation in reflectance that has been observed in experiments. At the specular backscatter, the reflectance α_1 is modeled with a $\mathcal{N}(1, 0.1)$ distribution. For the angled boards, the reflectance α_k is distributed as $\mathcal{N}(\alpha_0, 0.1\alpha_0)$. The first board is normal to the LOS. The other three boards are tilted at angles of $14^\circ, 31^\circ, 45^\circ$. A top-down view of the targets is shown in Fig. 6.1(a), and the intensity image as seen from the camera is shown in Fig. 6.1(b). The remaining simulation parameters are given in Table 6.1. A parametric waveform correlation routine selects optimal range and pulse width estimates, \hat{R} and $\hat{\sigma}_R$, respectively. Using the range estimates, the inclination angle at each point in the image is calculated, using 3×3 local windows and the plane-fitting algorithm as described in Section 5.4.1. Statistics of the inclination angle measurements are given in Table 6.2. For each object region in the image, the median inclination angle is selected as the estimator for the true inclination angle θ_i .

A reference surface is used as a calibration to calculate $\hat{\alpha}$ in each region. Having *a priori* knowledge of the image, it is simply a matter of selecting the pixels in the 0° angled board to use as a reference. Without the prior knowledge, or without a reference target, it would not be possible to verify the estimates experimentally. Estimates of the material reflectance require estimates of the effective gain and noise bias in the selected region and in a reference region, as well as an estimate of the inclination angle of the surface, resulting in $\hat{\alpha}_k$ having

Table 6.1. Simulation Waveform parameters.

| Parameter | Value |
|----------------|---|
| α | 0.7 |
| θ_i | $[0^\circ, 14^\circ, 31^\circ, 45^\circ]$ |
| σ_R | 2.7 ns |
| \mathfrak{G} | 0.2 mJ |
| B_n | 27400 |
| R_{min} | 156.7 m |

Table 6.2. Simulation. Inclination angle statistics. Angles are in degrees.

| Object | Actual | Median [$^\circ$] | Mean [$^\circ$] | Std Dev [$^\circ$] |
|--------------|--------|------------------------|----------------------|-------------------------|
| First Board | 0 | 0 | 0.9 | 5 |
| Second Board | 14 | 21 | 17.7 | 15.9 |
| Third Board | 31 | 28.2 | 34.8 | 13.8 |
| Fourth Board | 45 | 45 | 46.1 | 7.5 |

at least 5 degrees of freedom. The formula for the reflectance parameter estimator is

$$\hat{\alpha}_k = \frac{1}{\cos \hat{\theta}_k} \frac{\hat{G}_k - \hat{B}_n}{\hat{G}_{ref} - \hat{B}_n}, \quad k = 1, \dots, K \quad (6.1)$$

where the index k represents one of K segmented regions of the image corresponding to specific materials. Selected statistics of $\hat{\alpha}_k$ are shown in Table 6.3. To create the table data, the sample mean of the pixels in each segment is declared the effective gain estimate, and the estimated bias is subtracted from this value. In each region, the gain-bias difference is divided by the gain-bias difference of the reference segment, forming a ratio. The ratio of gain-bias differences is divided by the BRDF term of $\cos^2 \hat{\theta}_i$, which yields $\hat{\alpha}$.

6.2.2 Field Experiments.

The target setup consisted of flat foam and plywood boards, as well as a two-tone step target. The $20'' \times 30''$ foam board and the black-and-white step target were mounted vertically on posts, and the plywood board was leaned against the post at a 15° angle. The

Table 6.3. Simulation. Reflectance parameter statistics.

| Object | Actual | Median | Mean | Std Dev |
|--------------|--|--------|------|---------|
| First Board | $\alpha_n \sim \mathcal{N}(1, 0.1)$ | | 1 | |
| Second Board | $\alpha_2 \sim \mathcal{N}(\alpha_0, 0.1\alpha)$ | 0.73 | 0.74 | 0.11 |
| Third Board | $\alpha_3 \sim \mathcal{N}(\alpha_0, 0.1\alpha)$ | 0.64 | 0.65 | 0.09 |
| Fourth Board | $\alpha_4 \sim \mathcal{N}(\alpha_0, 0.1\alpha)$ | 0.70 | 0.71 | 0.10 |




Figure 6.1. Simulation target (a) range profile, (b) intensity profile




Figure 6.2. Peak amplitude of target boards with respect to aspect angle. Truth reflectance $\alpha = 0.7$.

two target setups are shown in Fig. 6.3. A 3° field-of-view lens on the ladar camera required the targets to be about 88 m distant in order to be in focus. A frame rate of 9 Hz was used to capture 144 frames of data per shot, and several shots of each target setup were taken. The camera sample rate was 434.5 MHz, resulting in anywhere from 43 to 60 samples, depending on the range gate settings. The multiple frames are combined, by sample, into a single frame, which was then cropped from 128×128 pixels down to 65×76 in order to emphasize the targets of interest. Assuming a Gaussian pulse shape, we use a correlation-maximization routine to model a waveform that closely matches the detected waveform in each pixel, which returns a range R and pulse width σ_w , which are the parameters of the waveform in the pixel; the pulse amplitude is just the peak value of the waveform.

Observing the laser intensity images of Fig. 6.3(c)-(d), we note that both the vertical black board and the tilted plywood board have reduced intensity compared to the vertical white board. By comparing the average waveform from the black and white board, in Fig. 6.4, it is clear that the loss of reflected energy from the black board reduces both the amplitude and pulse length of the received waveform, by almost 30%. To confirm this, we compute the bootstrapped mean of both the intensity and pulse width from each object in the image, with the results shown in Fig. 6.5. The statistics of the background are also included for comparison. If we assume the white board is normal to the incident laser pulse, then we can estimate the transmitted and reflected pulse width from that target to be ~ 3 ns. Because of the reduced reflectivity of the black board, the reflected pulse width from the vertical black board is measured nearly a full nanosecond shorter. Furthermore, since the plywood board is observed in Fig. 6.5(b) to be slightly more reflective than the black board, there is more energy in the reflected pulse; accordingly, the estimated value of the pulse width is 2.25 ns. The pulse width estimates corroborate neither intuition nor the cross-correlation method introduced in Chapter III.

The ranging results of the experiment are more satisfying. While the reflected pulse intensity from the black board and tilted board was reduced significantly, this did not affect the correlation routine for estimating the range to the targets. Fig. 6.6 shows the bootstrapped means of the range values for the targets in each setup. As expected, the

vertical boards are at the same range, 95.5m, while the tilted board is just slightly nearer. The background objects, which include the tree stump seen in Fig. 6.3, are distributed at a closer range, clearly distinct from the targets of interest.

Using sliding windows, we can calculate the local sample statistics at each (x, y) in the image. These values can also be bootstrapped, such that we can estimate the mean of the distribution of the local statistic. The local standard deviation window size was 5×5 , and the local skewness window was 7×7 . These local statistics were computed for both the intensity and estimated pulse width features of target setup 2. The bootstrapped means of each local statistic is shown in Fig. 6.7. In addition to the background, three target classes are observed: white board, black board, and tilted board. In all the cases, the mean of the background statistic is not equal to the means of the target statistics. However, the statistics of the tilted plywood board are similar to either the white or black board, depending on the local statistic being observed. The plywood and black board are similar in their bootstrapped mean value of the local standard deviation of intensity (in Fig. 6.7(a)), but not for local standard deviation of pulse width (in Fig. 6.7(c)).

The statistics presented in this section highlight the difficulty of accurately fitting waveforms to the experimental ladar data captured by the portable camera. In addition, the fitted waveform is of little analytic value if it is without a physical meaning. Additional experimentation and data analysis will be required to better quantify the observed measurement deficiencies.

6.3 Chapter Summary


This chapter described the feature selection procedure necessary for use in a classification routine. The feasibility of describing the distributions of objects in 3-D ladar data using nonparametric bootstrap resampling has been demonstrated by this experiment. Other nonparametric tests may also benefit the analysis, and shall be applied in the near future. Also, in order to draw better conclusions about modeling target surfaces for ladar image processing, this research requires data from a wider variety of targets at multiple ranges and orientations.



(a) Target setup 1



(b) Target setup 2



(c) Intensity image, setup 1



(d) Intensity image, setup 2

Figure 6.3. Color photos and intensity images of the two target setups. The plywood board in setup 2 is tilted at $\sim 15^\circ$ with respect to the vertical board.




Figure 6.4. Average waveforms reflected from the white and black boards, taken from Fig 6.3(c). The ratio of the peak intensity values is 0.62.

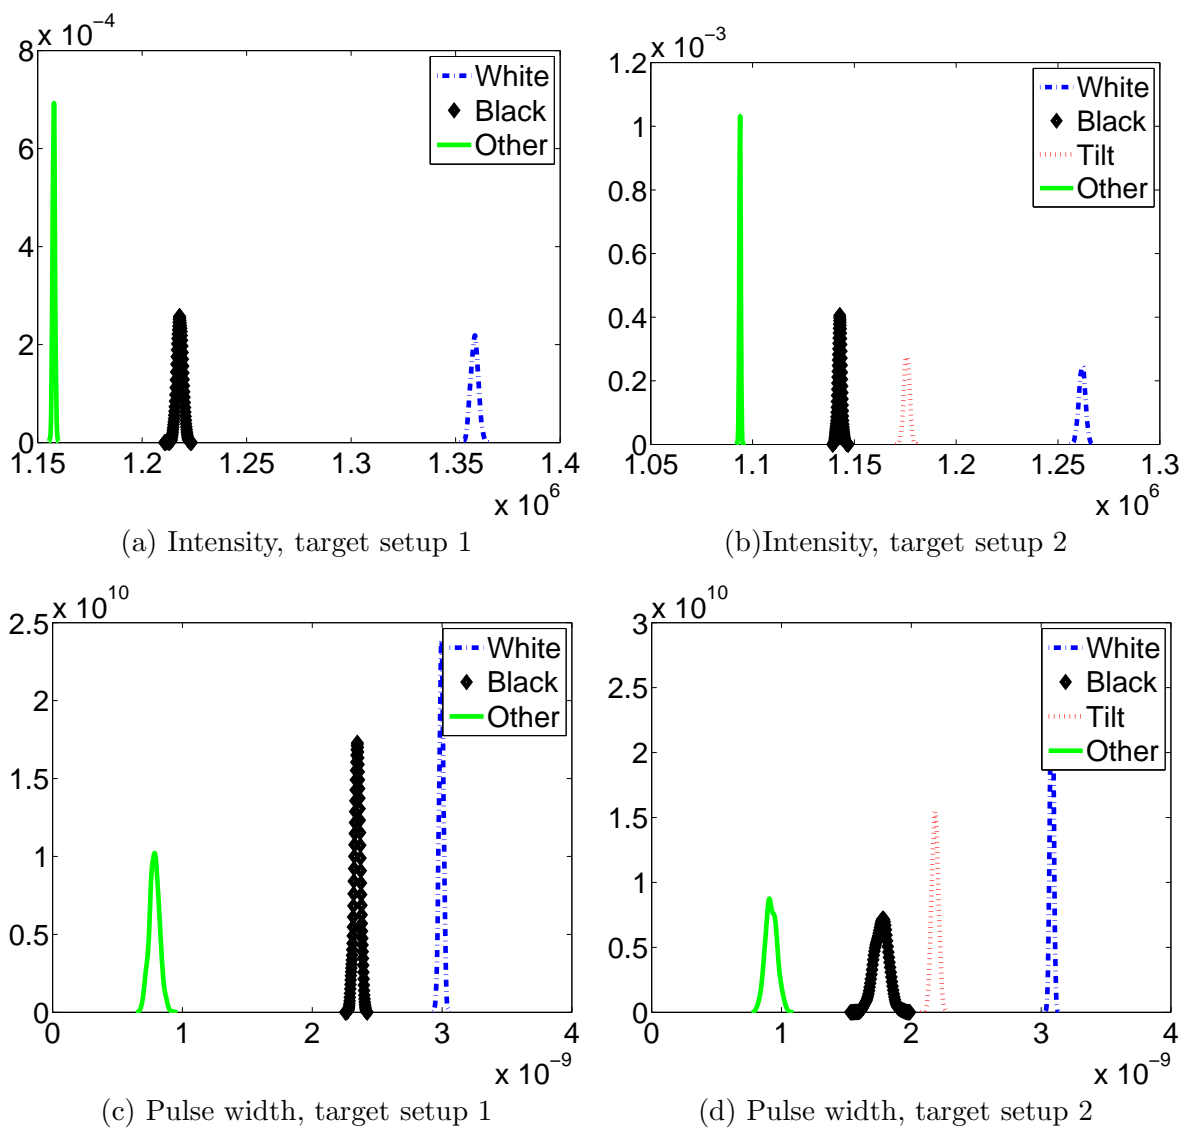


Figure 6.5. Bootstrapped mean value for parameters of hand-labeled objects from both target setups (1000 samples). Intensity is measured in photon counts, and pulse width is measured in seconds.




Figure 6.6. Bootstrapped mean range-to-targets.

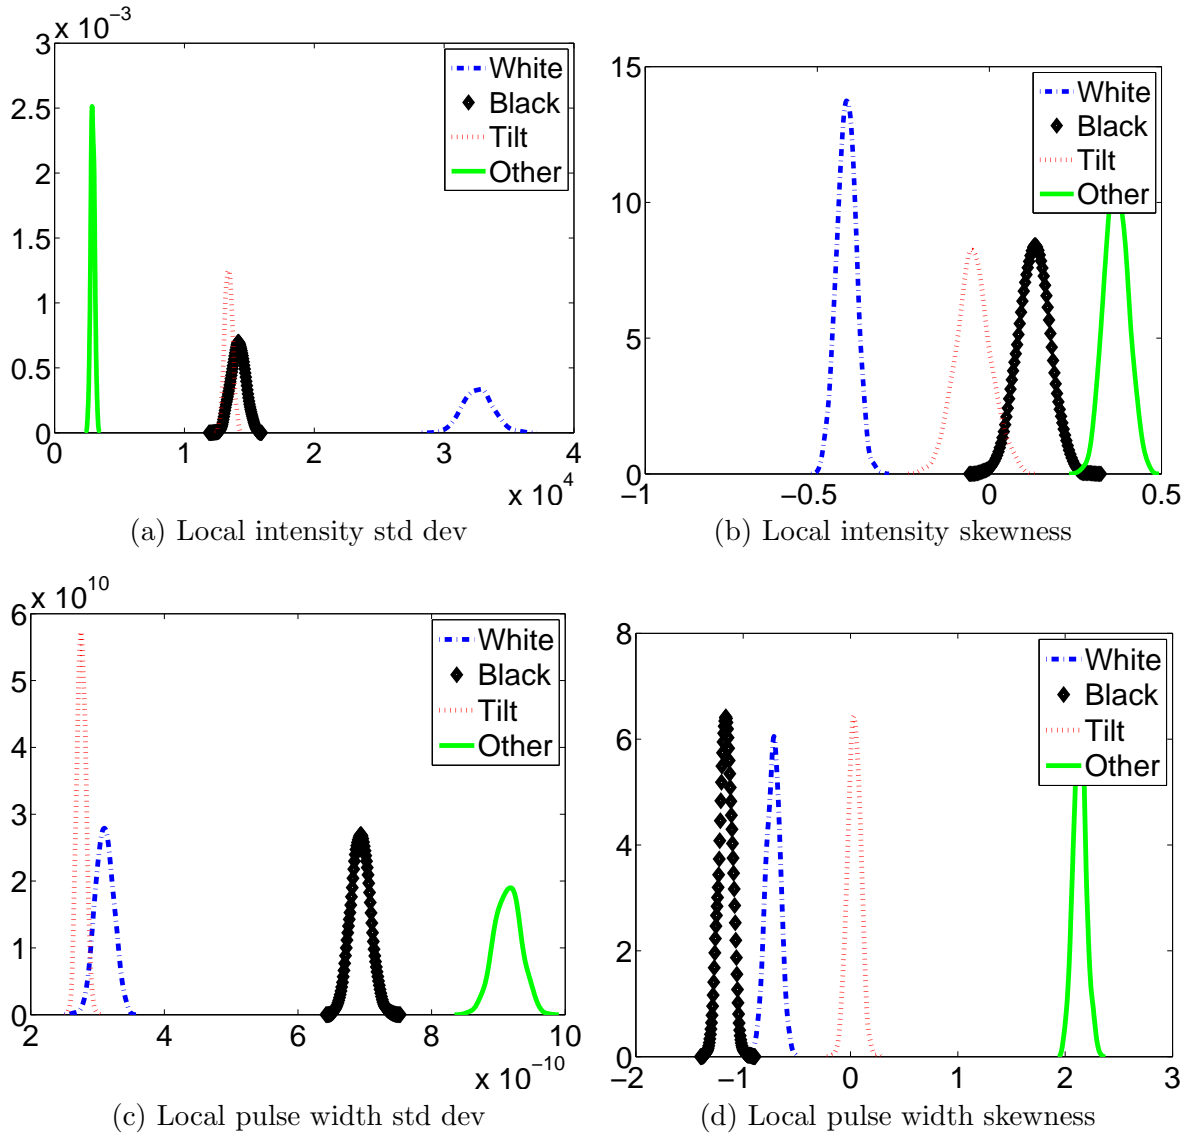


Figure 6.7. Bootstrapped means of local statistics (1000 samples). Standard deviation computed using 5×5 window, and skewness computed using 7×7 window.

VII. Nonparametric Classification for Ladar Images

Distribution-free classification methodology in image processing has been widely used in the medical and hyperspectral imaging communities in recent years [3] [105]. The multitude of approaches to the common issue of classification can generally be broken down into two types: supervised and unsupervised. The unsupervised, or data-driven, approach is preferred for this research, for reasons discussed in Chapter II. The contribution of this chapter is to combine the KDE-based range data segmentation method and the third-phase image generation of Chapter IV, and the feature analysis of Chapter VI into a usable classification scheme based on nonparametric tests. Using nonparametric tests is desirable for the classification, due to the irregular distribution of the parameters of interest: range, gain, inclination angle, etc. An effective nonparametric classification method currently in use is the k -nearest-neighbor (k NN) classifier [8] [104] [105]. The k NN rule classifies each of the regions in the unlabeled test by the majority label of its k -nearest neighbors in the training set.

The outline of the chapter is as follows. In Sec. 7.1 the nonparametric tests of location and dispersion are derived, and the likely classification features are discussed. The method is described in Sec. 7.2, and demonstrated in Sec. 7.3 on the data that has been used throughout the previous chapters. A summary of the method and notable contributions is provided in the final section.

7.1 Classification Tests and Features

In order to accomplish the relative classification of the segmented image regions, a combination of nonparametric tests are implemented. The segmented regions will be compared pairwise, resulting in a $R_\Sigma \times R_\Sigma$ matrix of outcomes for each hypothesis test. For pair of regions, the Wilcoxon rank-sum test and Ansari-Bradley test will be used to detect location and dispersion differences, respectively, between the sample data. In the context of ladar images, candidate data come from the segmented regions of any of the images described herein.

The combination of the location and dispersion tests for the two-sample situation has been suggested by [41]. Each method will test a hypothesis concerning the material reflectance in each segmented region. Combining the tests in this manner has a two-fold purpose. The location test poses the hypothesis that the median reflectance in each pair of regions is the same, versus the alternative that the medians are not equal. The objective of the test is to determine the effect of orientation and material on measured reflectance, that is, if changes in orientation (inclination angle) and target surface (or material) cause a shift in the median of the measured reflectance value. In a similar way, the dispersion test poses the scenario that the median of the two distributions are equal, but vary in dispersion (or variance). This test addresses the alternative to the location test, which is that the orientation and material of the illuminated objects affect the variation in reflectance, but not the median. The parametric corollary of the location test is the Student's t -test of equal means, and the corollary to the dispersion test is the F -test of equal variances.

The power of the nonparametric tests are such that any local statistic of the candidate data - mean, standard deviation, etc. - could also be used as the basis of the test. In the absence of absolute “ground truth” knowledge of the object parameters, the classification scheme will classify the regions in into “like,” “unlike, ” and “other” regions, building a table of the test outcomes between the regions.

In the following sections, the Wilcoxon and Ansari-Bradley tests are briefly described. For a more thorough discussion of these and other tests, the reader is directed to the excellent texts of [41] and [55].

7.1.1 Wilcoxon Rank Sum Test.

The Wilcoxon Rank Sum Test (WRST) tests the hypothesis that the two samples \mathbf{X} and \mathbf{Y} have the same probability distribution, and therefore the same location (or median). The only assumption is that \mathbf{X} and \mathbf{Y} are independent random samples from their respective distributions. Let $\mathbf{X} = [X_1, X_2, \dots, X_{n_1}]$ and $\mathbf{Y} = [Y_1, Y_2, \dots, Y_{n_2}]$ be two independent samples with unknown continuous density functions (CDFs) F_X and G_Y . The null hypothesis of the Wilcoxon test is that the continuous distribution function (CDFs)

are equal, against the alternative that they are unequal, that is

$$\begin{aligned} H_0 : F_X(x) &= G_Y(x), \forall x \\ H_A : F_X(x) &\neq G_Y(x), \text{ for some } x \end{aligned} \tag{7.1}$$

The test statistic W_n is calculated in the following manner. The $n = n_1 + n_2$ ranks of the combined sample $\mathbf{Z} = \{\mathbf{X}, \mathbf{Y}\}$ are determined. Then W_n is the sum of the ranks (1 to n) for the elements of \mathbf{X} in \mathbf{Z} , that is

$$W_n = \sum_{i=1}^n i S_i(\mathbf{Z}) \tag{7.2}$$

where $S_i(\mathbf{Z})$ is the indicator function defined as 1 if the i th ranked observation is from \mathbf{X} and as 0 if the observation is from \mathbf{Y} [55]. The null hypothesis is rejected if W_n is very large or very small, which would indicate that the samples vary significantly in location.

7.1.2 Ansari-Bradley Test.

The Ansari-Bradley test is a nonparametric test of dispersion used to test the hypothesis that two independent samples come from distributions with the same median. The construction of the test implies that if the null hypothesis is rejected (that is, the populations differ) it is because of unequal variance. Let $\mathbf{X} = [X_1, X_2, \dots, X_{n_1}]$ and $\mathbf{Y} = [Y_1, Y_2, \dots, Y_{n_2}]$ be two independent samples with unknown continuous density functions (CDFs) F_X and G_Y and a common median Δ . The $n = n_1 + n_2$ ranks of the combined sample $\mathbf{Z} = \{\mathbf{X}, \mathbf{Y}\}$ are determined. The test statistic C is the sum of the ranks for the observations of \mathbf{Y} in \mathbf{Z} , that is

$$C = \sum_{i=1}^n R_i \tag{7.3}$$

where R_i denotes the ranks of \mathbf{Y} . Finally, The ratio of the variance of the samples is given by the parameter $\gamma^2 = \text{var}(\mathbf{X})/\text{var}(\mathbf{Y})$. The two-tailed Ansari-Bradley test is given by

$$\begin{aligned} H_0 : \gamma^2 &= 1 \\ H_A : \gamma^2 &\neq 1 \end{aligned} \tag{7.4}$$

which is a test of the ratio of the variance from each sample. The null hypothesis is rejected for very large or small values of C , which would indicate the the samples have significant variation in dispersion.

7.1.3 Feature Selection.

Choosing appropriate feature data for the classification procedure is a critical task. The selection, therefore, is guided by examples provided throughout the literature. In the first case, the classification should be based on a generalized scheme [28]. This is to avoid overt specificity for particular classification problems. The second suggesting principal is to use local statistics of the feature data to guide classification [12] [60]. The method presented in this section satisfies both these criteria.

Implementing a generalized, data-driven approach to classification is necessary if a method is to be applied to disparate data sets. By using non-parametric tests for location and dispersion differences, the rank order of the data rather than the actual data value is used as the classification basis. It is the change in data from region to region that is therefore measured. Also, by using an initial range data segmentation, the method avoids any person-in-the-loop hand-labeling of the regions, a technique used in many other classification routines [97].

Concerning the use of local statistics, the concept of calculating properties (or saliencies) in a neighborhood for purposes of classification is well-known [37] [60] . In Sec. 4.3.1 the local range values were used to calculate a surface normal in each segmented region. For classification, however, raw range data is not appropriate, since the data is specific to the placement of each object in the range image. However, other ladar features, and the local statistics of those features, are appropriate for the classification.

The most promising features for classification are the system gain $\hat{\mathfrak{G}}_\alpha$ and the image intensity I . These features are related to each other, but the system gain estimate derives information from the incidence angle calculation, and is therefore expected to yield slightly better results. In addition to using these values in point form, the probability distributions of the local data statistics, calculated over moving windows, are also possible features for

Table 7.1. Classification Features

| Feature | Window Size |
|-------------------------------------|-------------|
| $\widehat{\mathfrak{G}}_\alpha$ | 1 |
| $sk(\widehat{\mathfrak{G}}_\alpha)$ | 8 |
| $var(I)$ | 15 |
| $ku(I)$ | 15 |

classification. In general it will not be necessary to use an elaborate feature extraction algorithm to reduce the number of available datasets, a situation that often appears in hyperspectral image analysis and other high-dimensional applications [105].

7.2 Classification Methodology

As has been stated previously in Sec. 7.1, the classifier assigns “like,” “unlike,” and “unknown” labels to the segmented regions of the ladar image, based on the selected features and local statistical measures within each region. The features are selected one at a time, and then each region is compared pairwise based on the probability distribution of that feature. All the tests are performed at the 10% significance level. A table of the features used in the nonparametric classification is given in Table. 7.1. With the exception of the first feature $\widehat{\mathfrak{G}}_\alpha$, all the entries in the table are local statistics of the second- and third-phase image data, calculated using the indicated neighborhood size. The local tests are, in order from top to bottom, the local skewness of $\widehat{\mathfrak{G}}_\alpha$, local variance of 2D intensity I , and the local kurtosis of I . The window sizes were determined using trial-and-error.

First, the location hypothesis (see (7.1)) is tested on a feature from Table. 7.1. For the demonstration of the method, the second feature, local $\widehat{\mathfrak{G}}_\alpha$ -skewness, calculated using an 8-pixel neighborhood, is used as the test feature. The segmented data is the ‘6 Feb 12’ Trial 4 dataset, which was segmented in Sec. 4.3.3; the segmentation is shown again in Fig. 7.2. The outcomes of the location tests are shown in Table. 7.2, where “1” rejects the null hypothesis, and “0” is a failure to reject, at 10% significance level. Histograms of the statistic in each region are shown in Fig. 7.1. The patch segmentation of the image, with each segmented numbered according to Table 7.2, is shown in Fig. 7.2(a).

According to segmented image in Fig. 7.2(a), regions 2, 3, and 4 correspond to the targets of interest, while the other regions may be assumed to be background, or other non-interesting targets. Looking at the corresponding rows in Table 7.2, the location hypothesis is rejected only in the 7th region, while the remaining regions fail to reject the null. The null hypothesis of equal median is then assumed, which is a necessary assumption for the Ansari-Bradley test of dispersion. The outcomes of the Ansari-Bradley dispersion test of the statistic are shown in Table 7.3. Entries in Table 7.3 which correspond to 1's (Reject H_0) in Table 7.2 are marked as “don't cares,” since the equal median assumption of the Ansari-Bradley test is violated. The remaining entries in Table 7.3 correspond to the Ansari-Bradley outcomes.

Again looking at the same three rows as before in Table 7.3, (and ignoring the “don't cares”), it appears that there is a pattern of “reject” and “fail to reject” the Ansari-Bradley test hypothesis in the rows of interest. Condensing the regions of interest to 2 – 4, (which correspond to the target boards and fence) and lumping the remaining regions into a generic “Other” category, the test results of Table 7.2 and Table 7.3 can be combined into a classification outcome table in Table 7.4. The table indicates similar location with L and similar dispersion (shape) with D . Regions 2 and 4 fail to reject for both tests, and so are labeled as similar. By contrast, Region 3 rejects the test of dispersion, so it is more similar to the “Other” regions.

But does this make sense from an object/target perspective? Consulting the segmented image in Fig. 7.2, the “similar” label, based on the local $\hat{\mathfrak{G}}_\alpha$ -skewness statistic, is applied to two of the prominent targets, while the angled target has a distribution more like that of the background. This result suggests that angled targets may be indistinguishable from background under certain statistical tests.

7.3 Examples

One more example is considered in this section, using another feature from Table 7.1: the local kurtosis of the $2D$ intensity I . Considering the same dataset (6 Feb 12 data, Trial 4), the location and dispersion tests are given in Table 7.5 and Table 7.6, respectively. The

Table 7.2. Wilcoxon Test Outcomes. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8

| | | Test Region | | | | | | | | |
|----------------|---|-------------|---|---|---|---|---|---|---|---|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Control Region | 1 | | 1 | 1 | 1 | 0 | 1 | 0 | 1 | 1 |
| | 2 | . | | 0 | 0 | 0 | 0 | 1 | 0 | 0 |
| | 3 | . | . | | 0 | 0 | 0 | 1 | 0 | 0 |
| | 4 | . | . | . | | 0 | 0 | 1 | 0 | 0 |
| | 5 | . | . | . | . | | 0 | 0 | 0 | 0 |
| | 6 | . | . | . | . | . | | 1 | 0 | 0 |
| | 7 | . | . | . | . | . | . | | 1 | 1 |
| | 8 | . | . | . | . | . | . | . | | 0 |
| | 9 | . | . | . | . | . | . | . | . | |

Table 7.3. Ansari-Bradley Test Outcomes. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8

| | | Test Region | | | | | | | | |
|----------------|---|-------------|---|---|---|---|---|---|---|---|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Control Region | 1 | | X | X | X | 1 | X | 1 | X | X |
| | 2 | . | | 1 | 0 | 0 | 0 | X | 0 | 1 |
| | 3 | . | . | | 1 | 1 | 1 | X | 1 | 0 |
| | 4 | . | . | . | | 0 | 0 | X | 0 | 1 |
| | 5 | . | . | . | . | | 0 | 0 | 0 | 1 |
| | 6 | . | . | . | . | . | | X | 0 | 1 |
| | 7 | . | . | . | . | . | . | | X | X |
| | 8 | . | . | . | . | . | . | . | | 1 |
| | 9 | . | . | . | . | . | . | . | . | |

Table 7.4. Classification Results. $sk(\hat{\mathfrak{G}}_\alpha)$, Window Size: 8

| | | Test Region | | | |
|----------------|-------|-------------|---|-----|-------|
| | | 2 | 3 | 4 | Other |
| Control Region | 2 | | L | L,D | L,D |
| | 3 | . | | L | L |
| | 4 | . | . | | L,D |
| | Other | . | . | . | |

final classification result, in Table 7.7 gives much the same result as the classification in Table 7.3, which is encouraging. In this case, however, the regions have a stronger grouping due to location of the statistic, rather than the dispersion.






Figure 7.1. Histograms of local statistic $sk(\hat{\Theta}_\alpha)$ for each segmented region. Window Size: 8 pixels. Dataset: 6 Feb '12, Trial 4.



(a) 6 Feb 12 Data, Trial 4 range segmentation



(b) 6 Feb 12 Data, Trial 4 System Gain Map

Figure 7.2. Patch segmentation and system gain map of 6 Feb 12 data, target 4.

Table 7.5. Wilcoxon Test Outcomes. $ku(I)$, Window Size: 15

| | | Test Region | | | | | | | | |
|----------------|---|-------------|---|---|---|---|---|---|---|---|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Control Region | 1 | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| | 2 | . | | 1 | 0 | 1 | 0 | 0 | 1 | 1 |
| | 3 | . | . | | 1 | 1 | 1 | 0 | 1 | 0 |
| | 4 | . | . | . | | 1 | 0 | 0 | 1 | 1 |
| | 5 | . | . | . | . | | 1 | 0 | 1 | 1 |
| | 6 | . | . | . | . | . | | 0 | 1 | 1 |
| | 7 | . | . | . | . | . | . | | 1 | 1 |
| | 8 | . | . | . | . | . | . | . | | 1 |
| | 9 | . | . | . | . | . | . | . | . | |

Table 7.6. Ansari-Bradley Test Outcomes. $ku(I)$, Window Size: 15

| | | Test Region | | | | | | | | |
|----------------|---|-------------|---|---|---|---|---|---|---|---|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Control Region | 1 | | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| | 2 | . | | X | 1 | X | 0 | 0 | X | X |
| | 3 | . | . | | X | X | X | 0 | X | 0 |
| | 4 | . | . | . | | X | 0 | 0 | X | X |
| | 5 | . | . | . | . | | X | 0 | X | X |
| | 6 | . | . | . | . | . | | 0 | X | X |
| | 7 | . | . | . | . | . | . | | X | X |
| | 8 | . | . | . | . | . | . | . | | X |
| | 9 | . | . | . | . | . | . | . | . | |

Table 7.7. Classification Results. $ku(I)$, Window Size: 15

| | | Test Region | | | |
|----------------|-------|-------------|---|---|-------|
| | | 2 | 3 | 4 | Other |
| Control Region | 2 | | - | L | L,D |
| | 3 | . | | - | L,D |
| | 4 | . | . | | L |
| | Other | . | . | . | |

7.4 Summary

A basic classification methodology based on nonparametric two-sample tests is described and implemented for use on segmented 3D ladar data. Nonparametric tests have been little used in the ladar imaging community, and this contribution presents classification techniques used in other fields (such as medical imaging) for consideration. Collective data points from each of the segmented regions are taken pairwise, and the Wilcoxon rank sum test is used to

test for location (median) of the distribution, and the Ansari-Bradley test is used to test for dispersion (variance). Necessary assumptions about the data are taken into consideration, and potential shortcomings of these assumptions, as well as the implementation of the tests themselves, are given a thorough treatment. Results of the classification methodology are mixed, but the technique is promising nonetheless, and comparisons are made with existing classifier models.

VIII. Conclusions

The contributions of this research increase the body of knowledge and the ultimate performance of three-dimensional flash lidar through a significantly updated waveform signal model, an analysis of the limitations of the pulse shape measuring capability of the lidar camera, and a novel, non-parametric based approach to object segmentation of lidar images.

8.1 Chapter Summaries

Chapter III detailed the procedure for experimental data collection using the ASC Portable Lidar camera, and reviewed the operating parameters of the camera. The initial data processing, including preparing the raw camera output data for the analysis software, is described. In addition, the compound-Poisson waveform model is proposed for use throughout the experiment. A widely-accepted ranging cross-correlation is described, and used as the basis for the contributions of Chapter IV.

Chapter IV introduced a novel lidar segmentation method. A short discussion of the nature of laser radar images was presented, and this motivated the necessity of classifying lidar images into distinct “phases” of imagery. The method is initialized using non-parametric probability density estimation, as this classifies the data into probability range bins, which is the first step in the segmentation. Plane fitting of each region is accomplished using principal components analysis, from which the plane inclination angle with respect to the lidar boresight is estimated. Also, the new third-phase image containing the relative material reflectance of the segmented object is estimated. A discussion of the strengths and weaknesses of the segmentation method was also presented. Comparison of the nonparametric segmentation with a traditional iterative seeded region growing approach shows that the two methods yield similar results, which is taken as validation of the new approach.

Chapter V presented an updated pulse model of a Poisson-distributed lidar signal. The amplitude of a Gaussian-shaped pulse model was shown to be dependent on the incidence angle of the incoming beam, as well as the optical reflectance of the illuminated material.

Parameter estimates for the range, gain, and signal bias were derived using a mixture of parametric and non-parametric matched filtering, as well as ML estimates. The RMSE of each parameter was evaluated using the CRB.

Chapter VI described the feasibility of finding the distributions of objects in 3D lidar data using nonparametric bootstrap resampling. Other nonparametric tests may also benefit the analysis, and suggestions are mentioned. Also, in order to draw better conclusions about modeling target surfaces for lidar image processing, this research requires data from a wider variety of targets at multiple ranges and orientations.

Chapter VII implemented a classification technique based on the outcomes of nonparametric hypothesis tests. Based on tests of location and dispersion, it is a unique approach to the unsupervised classification of lidar data images. Initial findings show that while not as accurate as other well known methods (such as the k NN classifier), it could be improved with a more structured approach to the classification criteria.

8.2 Summary of Contributions

8.2.1 KDE-Based Range Data Segmentation.

A novel solution for 3D lidar range data segmentation is the prime contribution of the research project. The segmentation is accomplished with a level of accuracy clearly superior to simple edge detection. Adaptation of a nonparametric kernel density estimation method is significant to the work; this technique, while common throughout the broad field of engineering, has seen little or no use within lidar applications. Segmentation via range data has enabled the remaining contributions, as it is the key to obtaining additional pieces of information about the scene.

8.2.2 Third-Phase Lidar Image Analysis.

In contrast with earlier scanning lidar systems, full-waveform flash lidar technology captures enormous amounts of data. The most basic visualization of this data is at the first-phase level, corresponding to principal range and intensity measurements; this type of

image is common throughout the literature. Successive derivations of the raw data lead to second-phase images (*i.e.* flat, 2D intensity images) and finally to third-phase images. This final type of image is synthesized from multiple waveforms captured by adjacent pixels, and conveys information that is not found easily in other ways, namely, object inclination angle and material reflectance. These new features form the core of both the new waveform model and the nonparametric classification scheme.

8.2.3 Backscatter Reflectance Waveform Model.

Measurement of waveform pulse expansion from resolved targets has been improved by integrating Lambertian reflection principles with the waveform model. Although consideration of backscatter reflectance is not a novel idea in itself, the approach for estimating aspect angle from pulse shape most certainly is. Leading researchers in this field have put forth many contributions in this area, and it would be beneficial to fuse some of those ideas and methods with more traditional image processing problems [48] [13]. By taking into account the effects of reflectance, the waveform model is improved over previous iterations, and better understanding of the underlying phenomena is achieved.

8.2.4 Image Classification using Nonparametric Tests.

An unusual contribution to the field of ladar image processing is made by demonstrating the feasibility of classification via nonparametric tests of location and dispersion. While this is common throughout other fields, especially medical imaging, it has not found great acceptance in the field at hand. It has been shown that while the classification method is workable, it depends heavily on the distributions of the features selected as classification variables. Emphasis has been placed on both properties of the material (such as reflectance) and on the distributions of the local statistics of those features.

8.3 Future Research Possibilities

There are numerous additional research opportunities that may be explored in the field of ladar image processing. Several of these areas are here described.

8.3.1 Integration of De-blurring and Range Data Segmentation.

The seminal work by McMahon [66] and the more recent work of Neff and Cain [68] [69] uses blind deconvolution to accomplish atmospheric de-blurring in 3D flash lidar data. While the present research has shown the importance of accounting for backscatter reflectance in the waveform model, it has made no attempt to simultaneously de-blur the image. Eliminating surface and range ambiguities while also performing range data segmentation would represent a valuable contribution toward the complete solution.

8.3.2 Anomaly Detection using 3D Flash Lidar.

The waveform correlation model put forth in [13] would benefit from a more rigorous analysis. Emphasis would be placed on unsupervised detection algorithms, drawing inspiration from the machine/computer vision fields. A savvy embellishment of some of the segmentation algorithms presented in this work, coupled with recently developed pulse shape measurement techniques, would be a welcome addition to this field.

8.3.3 Covariance of 3D Flash Lidar Systems.

The assumption of statistical independence of each pixel in the lidar detector array has been widely adopted [65], but not well studied. The local statistics of the waveform parameters calculated in [60] are clearly correlated with the neighboring pixels. Concerns about statistical independence in the data as well as the waveform model are on-going, although there is a dearth of this topic in the literature. Nonparametric and rank tests often can be found to give inconsistent results, due in part to violation of statistical independence which negates the form of the rejection regions of the test.

8.3.4 Full Waveform Analysis of Multiple Returns.

Lidar waveforms reflecting from foliage-covered areas typically have multiple pulse echoes or returns. The NOVAS cross-correlation algorithm put forth in [13], [24], is not currently well-suited to the multiple-echo situation, especially if the pulses overlap. Use

of the correlation algorithm could be expanded to include these scenarios, after being supplemented with rigorous waveform-fitting analysis. Further consideration of low-SNR waveforms, which occur especially at poorly-reflecting (such as tilted, or dispersive media) surfaces, should not be neglected.

Appendix A. Derivation of Cramer-Rao Bound

1.1 Preamble

The Poisson random variable $Y(t)$ is parameterized by the time varying function $\lambda(t)$, which is composed of both a signal and noise component:

$$\lambda(t|\mathbf{A}) = s(t, \mathbf{A}) + B_n \quad (\text{A.1})$$

where the signal part $s(t, \mathbf{A})$ is a known function of both t and \mathbf{A} , and B_n is a constant bias for the average noise power. Then, the probability that the number of events in the time interval (t_1, t_2) equals the integer k [72] is

$$P(N = k|\theta) = \frac{\left[\int_{t_1}^{t_2} \lambda(t|\theta) dt \right]^k}{k!} e^{-\int_{t_1}^{t_2} \lambda(t|\theta) dt} \quad (\text{A.2})$$

Note that for $k = 0$, (A.2) reduces to

$$P(N = 0|\theta) = \exp \left[- \int_{t_1}^{t_2} \lambda(t|\theta) dt \right] \quad (\text{A.3})$$

Suppose the interval $[0, T]$ is partitioned into M equally spaced intervals. Then $\tau = T/M$ is the length of each interval. Let $Q_i(\theta)$ be defined as a sampling of the rate function $\lambda(t|\theta)$ on the interval (a, b) , as in

$$Q_i(\theta) = \begin{cases} \lambda(t = b|\theta) \\ \frac{1}{b-a} \int_a^b \lambda(t|\theta) dt \end{cases} \quad (\text{A.4})$$

The first sampling function is simply the value of the rate function at time $t = b$, while the second function is the time average of $\lambda(t|\theta)$ over the interval (a, b) . We shall persist in using the generic function $Q_i(\theta)$ for now, although it will later become necessary to specify the sampling function. In light of this, we can modify (A.2) as follows. Let the time interval be given by $(a_i = \frac{T}{M}(i-1), b_i = \frac{iT}{M})$, $i = 0, 1, \dots, M$. Then, the probability of k_i events in

the interval (a, b) is given by

$$P(N_i = k_i | \theta) = \frac{Q_i(\theta)^{k_i}}{k_i!} e^{-Q_i(\theta)} \quad (\text{A.5})$$

The likelihood function for the set of M observations is simply the product of the probability in each time interval $\frac{T_i}{M}$, which we write as

$$P(\mathbf{N} = \mathbf{k} | \theta) = \prod_{i=1}^M \frac{Q_i(\theta)^{k_i}}{k_i!} e^{-Q_i(\theta)} \quad (\text{A.6})$$

The maximum likelihood estimate (MLE) for $\hat{\theta}$ is that value of θ that maximizes (A.6), or equivalently, maximizes the log of (A.6). The log-likelihood function is therefore

$$\begin{aligned} L(\theta | \mathbf{k}) &= \log P(\mathbf{N} = \mathbf{k} | \theta) \\ &= \sum_{i=1}^M k_i \log Q_i(\theta) - \sum_{i=1}^M Q_i(\theta) - \sum_{i=1}^M \log k_i! \end{aligned} \quad (\text{A.7})$$

We will assume $L(\theta | \mathbf{k})$, and therefore $\lambda(t | \theta)$, are differentiable in θ . Then, the possible MLE values $\hat{\theta}$ are those that maximize (A.7). In the following example we shall find the MLE for θ using both of the sampling functions defined in (A.4), and compare the variance of the estimate with the Cramer-Rao bound.

1.2 CRB

The unbiased estimators $\hat{\mathbf{A}} = [\hat{a}_1, \dots, \hat{a}_M]^T$ are bounded below by

$$\text{VAR}(\hat{a}_m - A_m) \geq (\mathbf{J})_{ii}^{-1} \quad (\text{A.8})$$

where J_{ii} is the ii th element in the $M \times M$ Fisher Information Matrix (FIM) \mathbf{J} . The elements of \mathbf{J} are:

$$\begin{aligned} J_{ij} &= E \left(\frac{\partial L_{\mathbf{Y}}(A)}{\partial A_i} \frac{\partial L_{\mathbf{Y}}(A)}{\partial A_j} \right) \\ &= -E \left(\frac{\partial^2 L_{\mathbf{Y}}(A)}{\partial A_i \partial A_j} \right) \end{aligned} \quad (\text{A.9})$$

Let $\hat{\mathbf{A}} = [R, G(\theta), B_n]^T$ be the 3×1 vector of parameter estimates. The FIM \mathbf{J} is a 3×3 matrix with diagonal elements given by

$$J_{ii} = -E \left(\frac{\partial^2 L_{\mathbf{Y}}(A)}{\partial^2 A_i} \right), i = 1, 2, 3 \quad (\text{A.10})$$

Evaluating (A.10) in two phases, the second partial derivatives of $L_{\mathbf{Y}}(A)$ are

$$\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial^2 A_i} = \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial^2 A_i} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{y_k}{\lambda^2(t, A)} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial A_i} \right)^2 \quad (\text{A.11})$$

Evaluating for each element of \mathbf{A} , the first derivatives are:

$$\begin{aligned} \frac{\partial L_{\mathbf{Y}}(\mathbf{A})}{\partial R} &= \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) \\ \frac{\partial L_{\mathbf{Y}}(\mathbf{A})}{\partial G} &= \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) \\ \frac{\partial L_{\mathbf{Y}}(\mathbf{A})}{\partial B} &= \sum_{k=1}^K \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right). \end{aligned} \quad (\text{A.12})$$

while the second partial derivatives are evaluated as

$$\begin{aligned} \frac{\partial^2 L_Y(A)}{\partial^2 R} &= \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial^2 R} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{y_k}{\lambda^2(t, A)} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial R} \right)^2 \\ \frac{\partial^2 L_Y(A)}{\partial^2 G} &= \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial^2 G} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{y_k}{\lambda^2(t, A)} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial G} \right)^2 \\ \frac{\partial^2 L_Y(A)}{\partial^2 B} &= - \sum_{k=1}^K \frac{y_k}{\lambda^2(t, A)} \end{aligned} \quad (\text{A.13})$$

The off-diagonal elements of \mathbf{J} require finding the mixed partial derivatives of $L_{\mathbf{Y}}(A)$, which

are:

$$\begin{aligned}
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial R \partial G} &= \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial R \partial G} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{y_k}{\lambda^2(t|\mathbf{A})} \\
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial R \partial B} &= \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial R \partial B} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{\partial \lambda(t|\mathbf{A})}{\partial B} \frac{y_k}{\lambda^2(t|\mathbf{A})} \\
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial G \partial R} &= \sum_{k=1}^K \frac{\partial^2 \lambda(t|\mathbf{A})}{\partial G \partial R} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{y_k}{\lambda^2(t|\mathbf{A})} \\
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial G \partial B} &= \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial G \partial B} \left(\frac{y_k}{\lambda(t|\mathbf{A})} - 1 \right) - \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{y_k}{\lambda^2(t|\mathbf{A})} \\
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial B \partial R} &= - \sum_{k=1}^K \frac{y_k}{\lambda^2(t|\mathbf{A})} \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \\
\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial B \partial G} &= - \sum_{k=1}^K \frac{y_k}{\lambda^2(t|\mathbf{A})} \frac{\partial \lambda(t|\mathbf{A})}{\partial G}
\end{aligned} \tag{A.14}$$

The calculation of the derivatives in (A.12) and (A.14) makes use of the fact that

$$\frac{\partial \lambda(t|\mathbf{A})}{\partial B} = 1. \tag{A.15}$$

Before evaluating the derivatives in (A.14), it is helpful to evaluate the negative expectation first (see (A.10)). Recall that the photocount event $Y(t)$ is a Poisson random variable with time varying parameter $\lambda(t|\mathbf{A})$, that is

$$Y(t) \sim \text{POI}(\lambda(t|\mathbf{A})) \tag{A.16}$$

Since the expectation (i.e. the mean) of a Poisson random variable returns this same parameter, then it follows that

$$E(Y(t)) = \lambda(t|\mathbf{A}). \tag{A.17}$$

This will aid the calculation of J. Proceeding, the negative expectations of the elements in

of (A.13) are

$$\begin{aligned}
J_{11} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial^2 R} \right) = \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial R} \right)^2 \\
J_{22} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial^2 G} \right) = \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial G} \right)^2 \\
J_{33} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial^2 B} \right) = \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})}
\end{aligned} \tag{A.18}$$

Note the similar form between the elements J_{11} and J_{22} . The remaining elements of \mathbf{J} are

$$\begin{aligned}
J_{12} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial R \partial G} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{1}{\lambda(t|\mathbf{A})} \\
J_{13} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial R \partial B} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{1}{\lambda(t|\mathbf{A})} \\
J_{21} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial G \partial R} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{\partial \lambda(t|\mathbf{A})}{\partial R} \frac{1}{\lambda(t|\mathbf{A})} \\
J_{23} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial G \partial B} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial G} \frac{1}{\lambda(t|\mathbf{A})} \\
J_{31} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial B \partial R} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial B} \frac{1}{\lambda(t|\mathbf{A})} \\
J_{32} &= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial B \partial G} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|\mathbf{A})}{\partial B} \frac{1}{\lambda(t|\mathbf{A})}
\end{aligned} \tag{A.19}$$

To complete the entries, it is only necessary to calculate the first partial derivatives of the waveform function $\lambda(t|\mathbf{A})$ with respect to each parameter in \mathbf{A} . The general form of the waveform function is given in (A.1), but the exact model is given by

$$\begin{aligned}
\lambda(t|\mathbf{A}) &= \frac{\mathfrak{G} \alpha \cos^2 \theta_i}{\sigma_R \sqrt{2\pi}} \exp \left\{ \frac{-1}{2\sigma_R^2} \left(t_k - \frac{2R}{c} \right)^2 \right\} + B_n \\
&= G \exp \left\{ \frac{-1}{2\sigma_R^2} \left(t_k - \frac{2R}{c} \right)^2 \right\} + B_n \\
&= s(t, \mathbf{A}) + B_n
\end{aligned} \tag{A.20}$$

In this model, the amplitude coefficient G is a function of the BRDF parameter $\alpha \cos^2 \theta_i$,

the pulse width σ_R , as well as the intrinsic gain value \mathfrak{G} itself. The functional dependence of these terms will be suppressed throughout the remaining notation.

With respect to R , the partial derivative of (A.20) is:

$$\begin{aligned}\frac{\partial \lambda(t|\mathbf{A})}{\partial R} &= \frac{2}{c} \frac{G}{\sigma_R} \left(t_k - \frac{2R}{c} \right) \exp \left\{ \frac{-1}{2\sigma_R^2} \left(t - \frac{2R}{c} \right)^2 \right\} \\ &= \frac{2}{c} \frac{1}{\sigma_R^2} \left(t_k - \frac{2R}{c} \right) \cdot s(t, \mathbf{A})\end{aligned}\tag{A.21}$$

where $s(t, \mathbf{A})$ is the signal part of the function as defined in (A.20). Moving on, the partial derivative with respect to effective gain G is

$$\begin{aligned}\frac{\partial \lambda(t|\mathbf{A})}{\partial G} &= \exp \left\{ \frac{-1}{2\sigma_R^2} \left(t_k - \frac{2R}{c} \right)^2 \right\} \\ &= \frac{1}{G} s(t, \mathbf{A})\end{aligned}\tag{A.22}$$

Now free to evaluate the elements of \mathbf{J} , we see that:

$$\begin{aligned}J_{11} &= -E \left(\frac{\partial^2 L_Y(A)}{\partial^2 R} \right) = \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})} \left(\frac{\partial \lambda(t|\mathbf{A})}{\partial R} \right)^2 \\ &= \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})} \left(\frac{2}{c} \frac{1}{\sigma_R^2} \left(t_k - \frac{2R}{c} \right) s(t_k, A) \right)^2 \\ &= \frac{4}{c^2} \frac{1}{\sigma_R^4} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right)^2 \frac{s^2(t_k, A)}{\lambda(t_k, A)} \\ &= \frac{4}{c^2} \frac{1}{\sigma_R^4} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right)^2 \frac{(\lambda(t_k, A) - B_n)^2}{\lambda(t_k, A)} \\ &= \frac{4}{c^2} \frac{1}{\sigma_R^4} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right)^2 \frac{(\lambda^2(t_k, A) - 2\lambda(t_k, A)B_n + B_n^2)}{\lambda(t_k, A)} \\ &= \frac{4}{c^2} \frac{1}{\sigma_R^4} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right)^2 \left(\lambda(t_k, A) - 2B_n + \frac{B_n^2}{\lambda(t_k, A)} \right)\end{aligned}\tag{A.23}$$

Continuing with the other diagonal elements of \mathbf{J} , we have

$$\begin{aligned}
J_{22} &= -E \left(\frac{\partial^2 L_Y(A)}{\partial^2 G} \right) \\
&= \sum_{k=1}^K \frac{1}{\lambda(t|A)} \left(\frac{\partial \lambda(t|A)}{\partial G} \right)^2 \\
&= \frac{1}{G^2} \sum_{k=1}^K \frac{s^2(t,A)}{\lambda(t|A)} \\
&= \frac{1}{G^2} \sum_{k=1}^K \frac{(\lambda(t|A) - B_n)^2}{\lambda(t|A)} \\
&= \frac{1}{G^2} \sum_{k=1}^K \left(\lambda(t|A) - 2B_n + \frac{B_n^2}{\lambda(t|A)} \right) \\
&= \frac{1}{G^2} \left[T_\lambda(A) - K2B_n + B_n^2 \sum_{k=1}^K \frac{1}{\lambda(t|A)} \right]
\end{aligned} \tag{A.24}$$

where the statistic $T_\lambda(A)$ is defined as

$$T_\lambda(A) = \sum_{k=1}^K \lambda(t|\mathbf{A}) \tag{A.25}$$

Finally, the last diagonal element of \mathbf{J} is

$$J_{33} = \sum_{k=1}^K \frac{1}{\lambda(t|\mathbf{A})} \tag{A.26}$$

Completing the off-diagonal elements of \mathbf{J} from (A.19), it follows that

$$\begin{aligned}
J_{12} &= J_{21} \\
&= -E \left(\frac{\partial^2 L_Y(\mathbf{A})}{\partial R \partial G} \right) \\
&= \sum_{k=1}^K \frac{\partial \lambda(t|A)}{\partial R} \frac{\partial \lambda(t|A)}{\partial G} \frac{1}{\lambda(t|A)} \\
&= \frac{2}{c} \frac{1}{\sigma_R^2} \frac{1}{G} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right) \frac{s^2(t,A)}{\lambda(t|A)} \\
&= \frac{2}{c} \frac{1}{\sigma_R^2} \frac{1}{G} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right) \left(\lambda(t|A) - 2B_n + \frac{B_n^2}{\lambda(t|A)} \right)
\end{aligned} \tag{A.27}$$

Also, the mixed partial derivatives with respect to range and signal bias,

$$\begin{aligned}
J_{13} &= J_{31} \\
&= -E \left(\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial R \partial B} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|A)}{\partial R} \frac{1}{\lambda(t|A)} \\
&= \frac{2}{c} \frac{1}{\sigma_R^2} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right) \frac{s(t,A)}{\lambda(t|A)} \\
&= \frac{2}{c} \frac{1}{\sigma_R^2} \sum_{k=1}^K \left(t_k - \frac{2R}{c} \right) \left(1 - \frac{B_n}{\lambda(t|A)} \right)
\end{aligned} \tag{A.28}$$

Finally, the mixed partial derivatives with respect to gain and signal bias,

$$\begin{aligned}
J_{23} &= J_{32} \\
&= -E \left(\frac{\partial^2 L_{\mathbf{Y}}(\mathbf{A})}{\partial G \partial B} \right) = \sum_{k=1}^K \frac{\partial \lambda(t|A)}{\partial G} \frac{1}{\lambda(t|A)} \\
&= \frac{1}{G} \sum_{k=1}^K \frac{s(t,A)}{\lambda(t|A)} \\
&= \frac{1}{G} \sum_{k=1}^K \left(1 - \frac{B_n}{\lambda(t|A)} \right)
\end{aligned} \tag{A.29}$$

The completed J matrix for the Poisson-noise signal of (A.2) is

$$\mathbf{J} = \begin{bmatrix} J_{11} & J_{12} & J_{13} \\ J_{21} & J_{22} & J_{23} \\ J_{31} & J_{32} & J_{33} \end{bmatrix} \tag{A.30}$$

where the elements of J are given in (A.23), (A.24), (A.26), (A.27), (A.28), and (A.29).

Appendix B. Derivation of Bayesian Abel Bound

2.1 Deterministic CRB

In which we present the background to evaluating the derivatives. We calculate the entries in the deterministic FIM, which require calculation of the second derivative of the log-likelihood function $L_{\mathbf{X}}(\boldsymbol{\theta})$. Taking the first derivative with respect to $\theta = R$, we have

$$\begin{aligned}\frac{\partial}{\partial \theta} L_X(\theta) &= -\frac{\partial}{\partial \theta} \frac{1}{2B} \sum_{k=1}^K (x_k - I_k(\theta))^2 \\ &= -\frac{1}{2B} \sum_{k=1}^K \left[-2 \frac{\partial I_k(\theta)}{\partial \theta} (x_k - I_k(\theta)) \right] \\ &= \frac{1}{B} \sum_{k=1}^K \left[x_k \frac{\partial I_k(\theta)}{\partial \theta} - I_k(\theta) \frac{\partial I_k(\theta)}{\partial \theta} \right]\end{aligned}\tag{A.1}$$

From this form, we take the second derivative with respect to θ :

$$\begin{aligned}\frac{\partial^2}{\partial \theta^2} L_X(\theta) &= \frac{\partial}{\partial \theta} \frac{1}{B} \sum_{k=1}^K \left[x_k \frac{\partial I_k(\theta)}{\partial \theta} - I_k(\theta) \frac{\partial I_k(\theta)}{\partial \theta} \right] \\ &= \frac{1}{B} \sum_{k=1}^K \left[\frac{\partial^2 I_k(\theta)}{\partial \theta^2} (x_k - I_k(\theta)) - \left(\frac{\partial I_k(\theta)}{\partial \theta} \right)^2 \right]\end{aligned}\tag{A.2}$$

Evaluating the derivatives of $I(t_k)$ we have the first derivative

$$\frac{\partial I(t_k)}{\partial \theta} = \frac{-4G(2\theta - t_k c)}{(c\sigma_w)^2} \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right]\tag{A.3}$$

and the second derivative

$$\frac{\partial^2 I(t_k)}{\partial \theta^2} = \frac{\partial}{\partial \theta} \frac{-4G(2\theta - t_k c)}{(c\sigma_w)^2} = \frac{-8G}{(c\sigma_w)^2} \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right]\tag{A.4}$$

Substituting (A.3) and (A.4) into (A.2), we have

$$\begin{aligned}\frac{\partial^2}{\partial \theta^2} L_X(\theta) &= \frac{1}{B} \sum_{k=1}^K \left[\frac{-8G}{(c\sigma_w)^2} (x_k - I_k(\theta)) - \left(\frac{-4G(2\theta - t_k c)}{(c\sigma_w)^2} \right)^2 \right] \\ &= -\frac{1}{B} \frac{8G}{(c\sigma_w)^2} \sum_{k=1}^K \left[(x_k - I_k(\theta)) + \frac{2G}{(c\sigma_w)^2} (2\theta - t_k c)^2 \right]\end{aligned}\quad (\text{A.5})$$

as the final form of the 2nd derivative of the log-likelihood function.

Since $X_k(\theta) \sim \mathcal{N}(I_k(\theta), B)$, it follows that $E[x_k(\theta)] = I_k(\theta)$. Now we take the negative expectation of (A.5) over the data x_k and the parameter θ ,

$$\begin{aligned}-E_{X_k} \left[\frac{\partial^2 L_X(\theta)}{\partial \theta^2} \right] &= \frac{1}{B} \frac{8G}{(c\sigma_w)^2} \sum_{k=1}^K E_{X_k} \left[(x_k - I_k(\theta)) + \frac{2G}{(c\sigma_w)^2} ((2\theta - t_k c))^2 \right] \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right] \\ &= \frac{1}{B} \frac{8G}{(c\sigma_w)^2} \sum_{k=1}^K \left[(E_{X_k}[x_k] - I_k(\theta)) + \frac{2G}{(c\sigma_w)^2} ((2\theta - t_k c))^2 \right] \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right] \\ &= \frac{1}{B} \left(\frac{4G}{(c\sigma_w)^2} \right)^2 \sum_{k=1}^K (2\theta - t_k c)^2 \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right]\end{aligned}\quad (\text{A.6})$$

This expectation can approximated as an integral over the rectangle function [45, p.71, eq. (3.25)], that is

$$-E_{X_k} \left[\frac{\partial^2 L_X(\theta)}{\partial \theta^2} \right] \approx f_s \frac{1}{B} \left(\frac{4G}{(c\sigma_w)^2} \right)^2 \int_{2\theta/c - \sigma_w}^{2\theta/c + \sigma_w} (\theta - t_k c)^2 dt \quad (\text{A.7})$$

Using a change of variable, $v = t_k c - 2\theta$, we can evaluate (A.7) as

$$\begin{aligned}-E_{X_k} \left[\frac{\partial^2 L_D(\theta)}{\partial \theta^2} \right] &\approx f_s \frac{1}{B} \left(\frac{4G}{(c\sigma_w)^2} \right)^2 \int_{-c\sigma_w}^{c\sigma_w} \frac{1}{c} (-v)^2 dv \\ &= f_s \frac{1}{B} \frac{32G^2}{3c^2\sigma_w} = \text{SNR} f_s \frac{32}{3c^2\sigma_w}\end{aligned}\quad (\text{A.8})$$

This approximation assumes the data was sampled at rate f_s .

2.2 Prior Information

The second derivative with respect to θ of the log of the prior distribution $p(\theta) \sim N(\mu_\theta, \sigma_\theta^2)$ is given by

$$\begin{aligned} \frac{\partial^2 \log p(\theta)}{\partial \theta^2} &= \frac{\partial}{\partial \theta} \left(\frac{\partial}{\partial \theta} \frac{(\theta - \mu_\theta)^2}{-2\sigma_\theta^2} \right) \\ &= \frac{1}{-2\sigma_\theta^2} \frac{\partial}{\partial \theta} (2(\theta - \mu_\theta)) \\ &= \frac{-1}{\sigma_\theta^2} \end{aligned} \tag{A.9}$$

Taking the negative expectation over the parameter θ simply yields

$$\begin{aligned} \mathbf{J}_P &= -E \left[\frac{-1}{\sigma_\theta^2} \right] \\ &= \frac{1}{\sigma_\theta^2} \end{aligned} \tag{A.10}$$

2.3 Bobrovsky-Zakai Bound

Van Trees [100] gives the form of the Bobrovsky-Zakai bound as

$$\zeta \geq \frac{h^2}{E_{x,\theta} \left[\left(\frac{p(\mathbf{x}, \theta+h)}{p(\mathbf{x}, \theta)} - 1 \right)^2 \right]} \tag{A.11}$$

This bound must be optimized over h , and converges to the Bayesian CRB in the limit as $h \rightarrow 0$.

The form of the Bobrovsky-Zakai bound for a Gaussian observation model is given in Appendix D of [81], and is repeated here.

$$BZB = \sup_h \frac{h^2}{\int_{\Theta} \frac{p^2(\theta+h)}{p(\theta)} e^{\frac{2}{\sigma^2} \|m(\theta+h) - m(\theta)\|^2} d\theta - 1} \tag{A.12}$$

The notation $\|\Phi\|^2$ indicates $\sum (\Phi \cdot \Phi^*)$, where $*$ denotes the complex conjugate. In the case of our particular model, $\mathbf{x} = \mathbf{I}(\theta) + \mathbf{n}$, and we will suppress the rectangle function for

the time being. The term $\|m(\theta + h) - m(\theta)\|^2$ can be written

$$\begin{aligned} \sum_{k=1}^K (I_k(\theta + h) - I_k(\theta))^2 &= \sum_{k=1}^K \left(G \left(1 - \left(\frac{2(\theta + h) - t_k c}{c\sigma_w} \right)^2 \right) - G \left(1 - \left(\frac{2\theta - t_k c}{c\sigma_w} \right)^2 \right) \right)^2 \\ &= \frac{16G^2 h^2}{(c\sigma_w)^4} \sum_{k=1}^K (2 + 2\theta - t_k c)^2 \end{aligned} \quad (\text{A.13})$$

We now recall that the pulse function $I_k(\theta)$ is multiplied by the rectangle function $\text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right]$. In (A.13), the summation term can be approximated as an integral over the duration of the rectangle function, as in

$$\sum_{k=1}^K (2 + 2\theta - t_k c)^2 \approx f_s \int_{2\theta/c - \sigma_w}^{2\theta/c + \sigma_w} (2 + 2\theta - t_k c)^2 dt = f_s \left(8\sigma_w + \frac{2c^2 \sigma_w^3}{3} \right) \quad (\text{A.14})$$

We substitute into (A.13) and simplify to obtain

$$\sum_{k=1}^K (I_k(\theta + h) - I_k(\theta))^2 \approx f_s \frac{32G^2 h^2 (12 + c^2 \sigma_w^2)}{3c^4 \sigma_w^3} \quad (\text{A.15})$$

which is independent of θ . The term $\int_{\Theta} p^2(\theta + h) / p(\theta) d\theta$, assuming a Gaussian *a priori* distribution on θ , becomes [31, 81]

$$\begin{aligned} \int_{\Theta} \frac{p^2(\theta + h)}{p(\theta)} d\theta &= \frac{1}{\sqrt{2\pi}\sigma_{\theta}} e^{\frac{2h^2 - 4h\mu + \mu^2}{-2\sigma_{\theta}^2}} \int_{\Theta} e^{\frac{\theta^2 + 2\theta(2h - \mu)}{-2\sigma_{\theta}^2}} d\theta \\ &= e^{\frac{2h^2 - 4h\mu + \mu^2}{-2\sigma_{\theta}^2} + \frac{(2h - \mu)^2}{2\sigma_{\theta}^2}} = e^{\frac{h^2}{\sigma_{\theta}^2}} \end{aligned} \quad (\text{A.16})$$

The Bobrovsky-Zakai bound is finally given by

$$BZB = \sup_h h^2 \left(e^{f_s \frac{64\text{SNR} h^2 (12 + c^2 \sigma_w^2)}{3c^4 \sigma_w^3} + \frac{h^2}{\sigma_{\theta}^2}} - 1 \right)^{-1} \quad (\text{A.17})$$

2.4 Bayesian Abel Bound

We take the expression for $\phi(h)$,

$$\begin{aligned}\phi(h) &= \frac{1}{h} \int_{\Theta} p_{\Theta}(\theta + h) \int_{\Omega} \frac{\partial \log p(\mathbf{x}|\theta)}{\partial \theta} p(\mathbf{x}|\theta + h) d\mathbf{x} d\theta \\ &\quad + \frac{1}{h} \int_{\Theta} \frac{\partial \log p_{\Theta}(\theta)}{\partial \theta} p_{\Theta}(\theta + h) d\theta\end{aligned}\tag{A.18}$$

The first term in (A.18) is given by

$$\begin{aligned}\int_{\Omega} \frac{\partial \log p(\mathbf{x}|\theta)}{\partial \theta} p(\mathbf{x}|\theta + h) d\mathbf{x} \\ &= -\frac{1}{2B} \int_{\Omega} \frac{\partial \|\mathbf{x} - \mathbf{I}(\theta)\|^2}{\partial \theta} p(\mathbf{x}|\theta + h) d\mathbf{x} \\ &= -\frac{1}{2B} \int_{\Omega} \frac{\partial (\mathbf{x} - \mathbf{I}(\theta))^T (\mathbf{x} - \mathbf{I}(\theta))}{\partial \theta} p(\mathbf{x}|\theta + h) d\mathbf{x}\end{aligned}\tag{A.19}$$

We use [99, eq. (A.371)] to simplify the integrand,

$$\begin{aligned}\frac{1}{B} \int_{\Omega} \frac{\partial \mathbf{I}^T(\theta)}{\partial \theta} (\mathbf{x} - \mathbf{I}(\theta)) p(\mathbf{x}|\theta + h) d\mathbf{x} &= \frac{1}{B} \frac{\partial \mathbf{I}^T(\theta)}{\partial \theta} \left(\int_{\Omega} \mathbf{x} p(\mathbf{x}|\theta + h) d\mathbf{x} - \mathbf{I}(\theta) \right) \\ &= \frac{1}{B} \frac{\partial \mathbf{I}^T(\theta)}{\partial \theta} (\mathbf{I}(\theta + h) - \mathbf{I}(\theta))\end{aligned}\tag{A.20}$$

The matrix multiplication effectively reduces to a summation over the time index t_k , yielding

$$\frac{4 \times \text{SNR}}{(c\sigma_w)^4} \sum_{k=0}^{N-1} (2\theta - ct_k) \left((2(\theta + h) - t_k c)^2 + (2\theta - t_k c)^2 \right) \text{rect} \left[\frac{2\theta - t_k c}{2c\sigma_w} \right]\tag{A.21}$$

As we have done previously, we approximate the summation over as an integral in time over the duration of the rectangle function.

$$\begin{aligned}&\approx f_s \frac{4 \times \text{SNR}}{(c\sigma_w)^4} \int_{2\theta/c - \sigma_w}^{2\theta/c + \sigma_w} (2\theta - ct_k) \left((2(\theta + h) - t_k c)^2 + (2\theta - t_k c)^2 \right) dt \\ &= f_s \frac{4 \times \text{SNR}}{(c\sigma_w)^4} \frac{1}{c} \int_{-c\sigma_w}^{c\sigma_w} (-u) \left((2h - u)^2 + (-u)^2 \right) du \\ &= f_s 4 \times \text{SNR} \left(\frac{8h^2}{3c^2\sigma_w} + \frac{\sigma_w}{5} \right)\end{aligned}\tag{A.22}$$

For the second term in (A.18), if we use the non-informative prior, $p(\theta) \sim \mathcal{U}(0, \theta_{max})$ the derivative vanishes. Therefore, we will use the prior Gaussian distribution:

$$\begin{aligned} \frac{1}{h} \int_{\Theta} \frac{\partial \log p_{\Theta}(\theta)}{\partial \theta} p_{\Theta}(\theta + h) d\theta &= -\frac{1}{h\sigma_{\theta}^2} \int_{\Theta} (\theta - \mu_{\theta}) p(\theta + h) d\theta \\ &= \frac{1}{\sigma_{\theta}^2} \end{aligned} \tag{A.23}$$

Finally, $\phi(h)$ is given by

$$\phi(h) = \frac{1}{\sigma_{\theta}^2} + f_s 4\text{SNR} \left(\frac{8h^2}{3c^2\sigma_w} + \frac{\sigma_w}{5} \right) \tag{A.24}$$

Bibliography

- [1] Alharthy, A. and J. Bethel. “Heuristic Filtering and 3D Feature Extraction From LIDAR Data”. *Int’l Archive of the Photogrammetry, Remote Sensing and Spatial Info Sciences*, 34(3A):29–34, 2002.
- [2] Anderson, T. K. “Kernel Density Estimation and K-means Clustering to Profile Road Accident Hotspots”. *Accident Analysis & Prevention*, 41(3):359–364, 2009.
- [3] Bandulasiri, A., R. N. Bhattacharya, and V. Patrangenaru. “Nonparametric inference for extrinsic means on size-and-(reflection)-shape manifolds with applications in medical imaging”. *Journal of Multivariate Analysis*, 100(9):1867–1882, 2009.
- [4] Bao, Y., C. Cao, C. Chang, X. Li, E. Chen, and Z. Li. “Segmentation To the Point Clouds of LIDAR Data Based on Change of Kurtosis”. *Proceedings of SPIE*, volume 6623, 66231N. 2007.
- [5] Bartels, M. and H. Wei. “Segmentation of LIDAR Data Using Measures of Distribution”. *Int’l Journal of Photogrammetry and Remote Sensing*, 36(7):426–431, 2006.
- [6] Bartels, M. and H. Wei. “Threshold-free object and ground point separation in LIDAR data”. *Pattern Recognition Letters*, 31(10):1089–1099, 2010.
- [7] Bartels, M., H. Wei, and D.C. Mason. “DTM Generation from LIDAR Data using Skewness Balancing”. *18th Int’l Conf on Pattern Recognition*, volume 1, 566–569. 2006.
- [8] Blanzieri, E. and F. Melgani. “Nearest Neighbor Classification of Remote Sensing Images with the Maximal Margin Principle”. *IEEE Trans. on Geoscience and Remote Sensing*, 46(6):1804–1811, June 2008.
- [9] Bobrovsky, B. Z. and M. Zakai. “A Lower Bound on the Estimation of Certain Diffusion Processes”. *IEEE Trans on Information Theory*, 22(1):45–52, Jan 1976.
- [10] Bors, A. G. and N. Nasios. “Kernel Density Estimation for Nonparametric Modeling”. *IEEE Trans on Systems, Man and Cybernetics*, 39(6):1543–1555, December 2009.
- [11] Bradley, A. P. “The Use of the Area Under the ROC Curve in the Evaluation of Machine Learning Algorithms”. *Pattern Recognition Letters*, 30(7):1145–1159, 1997.
- [12] Buschelman, E. A. and R. K. Martin. “Nonparametric Methods for Full-Waveform Ladar Images”. *IEEE Statistical Signal Processing Workshop*. Ann Arbor, MI, Aug 2012.
- [13] Cain, S. and B. Deas. “Anomaly detection using range profile and intensity signatures”. *Proceedings of SPIE*, volume 7687. 2010.
- [14] Cain S., R. Richmond and E. Armstrong. “Flash Light Detection and Ranging Accuracy Limits For Returns From Single Opaque Surfaces via Cramer-Rao Bounds”. *Applied Optics*, 45(24):6154–6162, August 2006.

- [15] Charaniya, A. P., R. Manduchi, and S. K. Lodha. "Supervised Parametric Classification of Aerial LiDAR Data". *Conference on Computer Vision and Pattern Recognition Workshop*, 30. 2004.
- [16] Chauve, A., C. Mallet, F. Bretar, S. Durrieu, P. Deseilligny, and W. Puech. "Processing Full-Waveform Lidar Data: Modelling Raw Signals". *ISPRS Workshop on Laser Scanning*. Espoo, Finland, Sep 2007.
- [17] Chehata, N., L. Guo, and C. Mallet. "Contribution of Airborne Full-Waveform Lidar and Image Data for Urban Scene Classification". *16th International Conference on Image Processing (ICIP)*, 1669–1672. Cairo, Egypt, Nov 2009.
- [18] Chen, Y., R. Yin, P. Flynn, and S. Broschat. "Aggressive Region Growing for Speckle Reduction In Ultrasound Images". *Pattern Recognition Letters*, 24:677–691, 2003.
- [19] Chu, C. C. and J.K. Aggarwal. "The Integration of Image Segmentation Maps Using Region and Edge Information". *IEEE Trans on Pattern Analysis and Machine Intelligence*, 15(12):1241 –1252, Dec 1993.
- [20] Chu, C. C., N. Nandhakumar, and J. K. Aggarwal. "Image Segmentation Using Laser Radar Data". *Pattern Recognition*, 23(6):569–581, 1990.
- [21] Clode, S., P. Kootsookos, and F. Rottensteiner. "The automatic extraction of roads from LIDAR data". *Int'l Journal of Photogrammetry and Remote Sensing*, 35:231–237, 2004.
- [22] Clode, S., F. Rottensteiner, P. Kootsookos, and E.E. Zelniker. "Detection and vectorization of roads from lidar data". *Photogrammetric Engineering and Remote Sensing*, 73(5):517–535, 2007.
- [23] Cortes, C. and M. Mohri. "Confidence intervals for the area under the ROC curve". *Advances in Neural Information Processing Systems*, 17:305–312, 2004.
- [24] Deas, B. T. *Pulse Shape Correlation for Laser Detection and Ranging*. Masters thesis, Air Force Institute of Technology, WPAFB, OH, Sept 2010.
- [25] Dima, C., M. Hebert, and A. Stentz. "Enabling learning from large datasets: Applying active learning to mobile robotics". *Proc of the Intl Conf on Robotics and Automation*, 1:108–114, May 2004.
- [26] Elgammal, A., R. Duraiswami, D. Harwood, and L. S. Davis. "Background and Foreground Modeling Using Nonparametric Kernel Density Estimation for Visual Surveillance". *Proceedings of IEEE*, 90(7):1151–1163, 2002.
- [27] Fawcett, T. "An introduction to ROC Analysis". *Pattern Recognition Letters*, 27:861–874, 2006.
- [28] Gao, Y., A. Rehman, and Z. Wang. "CW-SSIM Based Image Classification". *IEEE Int'l Conf. on Image Proc (ICIP)*. Brussels, Belgium, 2011.
- [29] Gonzalez, R. C., R. E. Woods, and S. L. Eddins. *Digital Image Processing Using MATLAB*. Pearson Prentice Hall, New Jersey, 2004.

- [30] Goodman, J. W. *Statistical Optics*. John Wiley and Sons, New York, New York, 1985.
- [31] Gradshteyn, I.S. and I.M. Ryzhik. *Table of Integrals, Series and Products*. Academic Press, Inc., San Diego, CA, 5th edition, 1994.
- [32] Grandell, J. *Mixed Poisson Processes*. Chapman & Hall, London, 1997.
- [33] Greer, D.R., I. Fung, and J.H. Shapiro. “Maximum-likelihood multiresolution laser radar range imaging”. *IEEE Trans on Image Processing*, 6(1):36–46, 2002.
- [34] Gross, H., B. Jutzi, and U. Thoennessen. “Segmentation of tree regions using data of a full-waveform laser”. *Int’l Journal of Photogrammetry and Remote Sensing*, 36(part 3):57–62, 2007.
- [35] Gross, H., B. Jutzi, and U. Thoennessen. “Intensity Normalization by Incidence Angle and Range of Full-Waveform Lidar Data”. *Proc. of Commission IV, Int’l Archive of Photogrammetry, Remote Sensing, and Spatail Info. Sci.*, 37:405–412, 2008.
- [36] Hancock, J., M. Hebert, and C. Thorpe. “Laser-Intensity Based Obstacle Detection”. *Proc of the 1998 IEEE/RSJ Intl Conf on Intelligent Robots and Systems*. Victoria, B.C., Canada, Oct. 1998.
- [37] Hebert, M. and N. Vandapel. “Terrain Classification Techniques From LADAR Data for Autuonomous Navigation”. *Robotics Institute*, 411, 2003.
- [38] Hegde, G. M. and C. Ye. “A Recursive Planar Feature Extraction Method for 3D Range Data Segmentation”. *IEEE Int’l Conf on Systems, Man, and Cybernetics*, 3119–3124. Oct 2011.
- [39] Hernandez-Marin, S., A.M. Wallace, and G.J. Gibson. “Bayesian analysis of lidar signals with multiple returns”. *IEEE Trans on Pattern Analysis and Machine Inteligence*, 2170–2180, 2007.
- [40] Hoffman, R. and A. K. Jain. “Segmentation and Classification of Range images”. *IEEE Trans. on Pattern Analysis and Machine Intel.*, PAMI-9(5):608–620, Sep 1987.
- [41] Hollander, M. and D. A. Wolfe. *Nonparametric Statistical Methods*. John Wiley & Sons, New York, NY, 1973.
- [42] Huang, J., A.B. Lee, and D. Mumford. “Statistics of Range Images”. *IEEE Conf on Computer Vision and Pattern Recognition*, volume 1, 324–331. Jun 2000.
- [43] I.Reuven and H. Messer. “A Barankin-Type Lower Bound on the Estimation Error of a Hybrid Parameter Vector”. *IEEE Transasctions on Information Theory*, 43(3):1084–1093, May 1997.
- [44] Jain, A., T. Newman, and M. Goulish. “Range-Intensity Histogram for Segmenting LADAR Images”. *Pattern Recognition Letters*, 13(1):41–56, 1992.
- [45] Johnson, S. *Range Precision of LADAR Systems*. Doctoral dissertation, Air Force Institute of Technology, WPAFB, OH, Sept 2008.

- [46] Johnson, S. E. and S. C. Cain. “Bound on Range Precision for Shot-Noise Limited Ladar Systems”. *Applied Optics*, 47(28):5147–5154, Oct 2008.
- [47] Jordan, S. “Range Estimation Algorithms Comparison in Simulated 3-D Flash LADAR Data”. *IEEE Aerospace Conference*, 1–7. March 2009.
- [48] Jutzi, B., B. Eberle, and U. Stilla. “Estimation and measurement of backscattered signals from pulsed laser radar”. *Proceedings of SPIE*, volume 4885, 256–267. 2003.
- [49] Jutzi, B. and U. Stilla. “Measuring and processing the waveform of laser pulses”. *Optical*, 194–203, 2005.
- [50] Jutzi, B. and U. Stilla. “Range Determination with waveform recording laser systems using a Wiener Filter”. *Int’l Journal of Photogrammetry and Remote Sensing*, 61(2):95–107, 2006.
- [51] Kay, S. M. *Fundamentals of Statistical Signal Processing*. Prentice Hall, Upper Saddle River, New Jersey, 1998.
- [52] Kervrann, C. and F. Heitz. “A Markov Random Field Model-based Approach to Unsupervised Texture Segmentation Using Local and Global Spatial Statistics”. *IEEE Transactions on Image Processing*, 4(6):856–862, 6 1995.
- [53] Khotanzad, A. and J.-Y. Chen. “Unsupervised segmentation of textured images by edge detection in multidimensional feature”. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 11(4):414–421, April 1989.
- [54] Kirchhof, M., B. Jutzi, and U. Stilla. “Iterative processing of laser scanning data by full waveform analysis”. *Int’l Journal of Photogrammetry and Remote Sensing*, 63(1):99–114, 2008.
- [55] Kvam, P. H. and B. Vidakovic. *Nonparametric Statistics with Applications to Science and Engineering*. John Wiley & Sons, Hoboken, NJ, 2007.
- [56] Liu, Jingang, Deyuan Shen, Siu-Chung Tam, and Yee-Loy Lam. “Modeling pulse shape of Q-switched lasers”. *IEEE Journal of Quantum Electronics*, 37(7):888–896, Jul 2001.
- [57] Liu, X. and D. Wang. “Image and Texture Segmentation Using Local Spectral Histograms”. *IEEE Trans on Image Processing*, 15(10):3066–3077, 2006.
- [58] Lu, D. and Q. Weng. “A survey of image classification methods and techniques for improving classification performance”. *International Journal of Remote Sensing*, 28(5):823–870, 2007.
- [59] Luo, W. “An Efficient Detail-Preserving Approach for Removing Impulse Noise in Images”. *IEEE Signal Proc. Letters*, 13(7):413–416, July 2006.
- [60] Macedo, J., R. Manduchi, and L. Matthies. “Ladar-based Discrimination of Grass From Obstacles for Autonomous Navigation”. *Intl Symposium on Experimental Robotics*. Honolulu, HI, Oct 2000.

- [61] Maini, R. and H. Aggarwal. “Study and Comparison of Various Image Edge Detection Techniques”. *Int’l Journal of Image Proc.*, 3(1):1, 2009.
- [62] Mallet, C., U. Soergel, and F. Bretar. “Analysis of Full-Waveform Data for Classification of Urban Areas”. *Photogrammetrie Fernerkundung Geoinformation*, 5:337–349, 2008.
- [63] Masson, P. and W. Pieczynski. “SEM Algorithm and Unsupervised Statistical Segmentation of Satellite Images”. *IEEE Trans. on Geoscience and Remote Sensing*, 31(3):618–633, 1993.
- [64] McAulay, R. and L. Seidman. “A useful form of the Barankin lower bound and its application to PPM threshold analysis”. *IEEE Trans on Information Theory*, 15(2):273–279, 1969.
- [65] McMahon, J. R. *Improving Range Estimation of a 3-Dimensional FLASH LADAR via Blind Deconvolution*. Doctoral dissertation, Air Force Institute of Technology, WPAFB, OH, Sept 2010.
- [66] McMahon, J.R., R.K. Martin, and S.C. Cain. “Three-dimensional FLASH laser radar range estimation via blind deconvolution”. *Journal of Applied Remote Sensing*, 4:043517, 2010.
- [67] Nakas, C. T. and C. T. Yiannoutsos. “Ordered multiple-class ROC analysis with continuous measurements”. *Statistics in Medicine*, 23(0R):3437–3449, 2004.
- [68] Neff, B. J. and S. C. Cain. “Multiple Surface Discrimination in Three-Dimensional FLASH Laser Radar While Minimizing the Effects of Diffraction”. *Optical Engineering*, 51(5):056201, 2012.
- [69] Neff, B. J., S. C. Cain, and R. K. Martin. “Discrimination of Multiple Ranges Per Pixel in 3D Flash LADAR While Minimizing the Effects of Diffraction”. *Proceedings of SPIE*, volume 8165, 81050J. 2011.
- [70] Nitzan, D., A. Brain, and R. Duda. “The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis”. *Proceedings of IEEE*, 65(2):206–220, February 1977.
- [71] Pal, N.R. and S.K. Pal. “A review on image segmentation techniques”. *Pattern Recognition*, 26(9):1277–1294, 1993.
- [72] Papoulis, A. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, Inc., 3rd. edition, 1991.
- [73] Persson, A. *Extraction of Individual Trees Using Laser Radar Data*. Technical Report FO-R-0236-SE, Swedish Defence Research Agency, 2001.
- [74] Persson, A., U. Sderman, J. Tpel, and S. Ahlberg. “Visualization and analysis of full-waveform airborne laser scanner data”. *Int’l Journal of Photogrammetry and Remote Sensing*, 36(part 3):103–108, 2005.

- [75] Peterson, C. J. R. *Near Earth Object Detection Using a Poisson Statistical Model for Detection on Images Modeled from the Panoramic Survey Telescope & Rapid Response System*. Masters thesis, Air Force Institute of Technology, WPAFB, OH, March 2012.
- [76] Quelle, H.C., Y. Delignon, and A. Marzouki. “Unsupervised Bayesian segmentation of SAR images using the Pearson system distributions”. *Int’l Geoscience and Remote Sensing Symposium*, 1538–1540. 2002.
- [77] Reece, S. and D. Nicholson. “Tighter alternatives to the Cramer-Rao lower bound for discrete-time filtering”. *8th Int’l Conf on Information Fusion*, volume 1, 6 pp. July 2005.
- [78] Reitberger, J., P. Krzystek, and U. Stilla. “Benefit of Airborne Full Waveform LIDAR for 3D segmentation and classification of single trees”. *Annual Conf. of American Society for Photogrammetry and Remote Sensing*. Mar 2009.
- [79] Reitberger, J., C. Schnrr, P. Krzystek, and U. Stilla. “3D segmentation of single trees exploiting full waveform LIDAR data”. *Int’l Journal of Photogrammetry and Remote Sensing*, 64(6):561–574, 2009.
- [80] Renaux, A., P. Forster, P. Larzabal, and C. Richmond. “The Bayesian Abel Bound on the Mean Square Error”. *Int’l Conf on Acoustics, Speech and Signal Proc.*, 9–12. May 2006.
- [81] Renaux, A., P. Forter, P. Larzabal, C.D. Richmond, and A. Nehorai. “A Fresh Look at the Bayesian Bounds of the Weiss-Weinstein Family”. *IEEE Trans on Signal Processing*, 56(11):5344–5352, Nov 2008.
- [82] Richmond, R. D. and S. C. Cain. *Direct-Detection LADAR Systems*. SPIE Press, Bellingham, WA, 2010.
- [83] Robinson, D. and P. Milanfar. “Fundamental Performance Limits in Image Registration”. *IEEE Transactions on Image Processing*, 13(9):1185–1199, Sept 2004.
- [84] Samadzadegan, F., M. Hahn, and B. Bigdeli. “Automatic road extraction from LIDAR data based on classifier fusion”. *Joint Urban Remote Sensing Event*, 1–6. May 2009.
- [85] Sampat, M.P., A.C. Patel, Y. Wang, S. Gupta, C.W. Kan, A.C. Bovik, and M.K. Markey. “Indexes for Three-Class Classification Performance AssessmentAn Empirical Comparison”. *IEEE Trans on Information Technology in Biomedicine*, 13(3):300–312, 2009.
- [86] Schott, J. R. *Fundamentals of Polarimetric Remote Sensing*. SPIE Press, Bellingham, WA, 2009.
- [87] Schubert, C. M., S. N. Thorsen, and M. E. Oxley. “The ROC Manifold for Classification Systems”. *Pattern Recognition*, 44(2):350–362, Feb 2011.
- [88] Seidman, L.P. “Performance limitations and error calculations for parameter estimation”. *Proceedings of the IEEE*, 58(5):644–652, 1970.

- [89] Shapiro, J.H. “Bounds on the Area Under the ROC curve”. *Journal of the Optical Society of America*, 16(1):53–57, 1999.
- [90] Sok, C. and M. Adams. “Visually Aided Feature Extraction from 3D Range Data”. *Proc. of the IEEE Int’l Conf. on Robotics and Automation*, 2273–2279. Anchorage, AK, May 2010.
- [91] Song, J.H., S.H. Han, KY Yu, and Y.I. Kim. “Assessing the possibility of land-cover classification using lidar intensity data”. *Int’l Journal of Photogrammetry and Remote Sensing*, 34(3/B):259–262, 2002.
- [92] Steinvall, O. “Effects of target shape are reflection on laser radar cross sections”. *Applied Optics*, 39(24):4381–4391, Aug 2000.
- [93] Stettner, R., H. Bailey, R. Richmond, and Inc. Advanced Scientific Concepts. “Eye-safe laser radar 3-D imaging”, 2005.
- [94] Stettner, R., H. Bailey, and S. Silverman. “Three Dimensional Flash LADAR Focal Planes and Time Dependent Imaging”. *Int’l Journal of High Speed Electronics and Systems*, 18(2):401–406, 2008.
- [95] Taylor, G. and L. Kleeman. “Robust Range Data Segmentation Using Geometric Primitives for Robotic Applications”. *Proc. of 9th Int’l Conf on Signal and Image Proc.* 2003.
- [96] Tsai, A., J. Zhang, and A.S. Willskey. “Expectation-maximization algorithms for image processing using multiscale models and mean- field theory, with applications to laser radar range profiling and segmentation”. *Optical Engineering*, 40(7):1287–1301, Jul 2001.
- [97] Vailaya, A., A. Jain, and H. J. Zhang. “On Image Classification: City Images vs. Landscapes”. *Pattern Recognition*, 31(12):1921–1935, 1998.
- [98] Van Trees, H. L. *Detection, Estimation, and Modulation Theory*, volume I. John Wiley and Sons, Inc., New York, New York, 1968.
- [99] Van Trees, H. L. *Detection, Estimation, and Modulation Theory, Part IV: Optimum Array Processing*, volume IV. John Wiley and Sons, Inc., New York, NY, 2002.
- [100] Van Trees, H.L. and K.L. Bell (editors). *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*. IEEE Press, 2007.
- [101] Vandapel, N., D. F. Huber, A Kapuria, and M. Hebert. “Natural Terrain Classification Using 3-D LADAR Data”. *Int’l Conf on Robotics and Automation*, volume 5, 5117–5122. 2004.
- [102] Vasile, A. and R. M. Marino. “Pose-Independent Automatic Target Detection and Recognition using 3-D Ladar Data”. *Proceedings of SPIE*, 5426:67–83, 2004.
- [103] Wagner, W., A. Roncat, T. Melzer, and A. Ullrich. “Waveform analysis techniques in airborne laser scanning”. *ISPRS Workshop on Laser Scanning*, 3:413–418, Sep 2007.

- [104] Weinberger, K. Q. and L. K. Saul. “Distance Metric Learning for Large Margin Nearest Neighbor Classification”. *Journal of Machine Learning Research*, 10:207–244, June 2009.
- [105] Yang, J., P. Yu, and B. Kuo. “A Nonparametric Feature Extraction and Its Application to Nearest Neighbor Classification for Hyperspectral Image Data”. *IEEE Trans. on Geoscience and Remote Sensing*, 48(3):1279–1293, March 2010.
- [106] Yao, W., S. Hinz, and U. Stilla. “Extraction and Motion Estimation of Vehicles in single-pass Airborne LiDAR Data Towards Urban Traffic Analysis”. *Int’l Journal of Photogrammetry and Remote Sensing*, 66(3):260–271, 2011.
- [107] Ye, J., A. Wallace, and J. Thompson. “Parallel Markov Chain Monte Carlo Computation for Varying-Dimension Signal Analysis”. *17th European Signal Processing Conference*, 2673 –2677. Glasgow, Scotland, Aug 2009.
- [108] Yunfei, B., L. Guopingb, C. Chunxiang, L. Xiaowen, Z. Haoa, H. Qishenga, B. Linyana, and C. Chaoyia. “Classification of Lidar Point Cloud and Generation of DTM from LiDAR Height and Intensity Data in Forested Area”. *Int’l Journal of Photogrammetry and Remote Sensing Congress*. 2008.
- [109] Zweig, M. H. and G. Campbell. “Reciever-Operating Characteristic (ROC) Plots: A Fundamental Tool in Clinical Medicine”. *Clinical Chemistry*, 39(4):561–577, 1993.

Vita

Capt Eric A. Buschelman is a 2005 Distinguished Graduate of the USAF ROTC program at the University of Kansas, where he earned the Bachelor of Science in Electrical Engineering and received a commission in the United States Air Force. Capt Buschelman's first assignment was to the Air Force Institute of Technology at Wright-Patterson AFB, OH where he received the Masters of Science in Electrical Engineering in March 2007.

Capt Buschelman was then assigned to the National Air and Space Intelligence Center at Wright-Patterson AFB as a Foreign Airborne Radar Analyst. While there, he attended the Air and Space Basic Course in Maxwell AFB, AL, where he was chosen as Outstanding Contributor to the flight by his peers. In March 2009, Capt Buschelman was selected for position of Deputy Senior Duty Officer of the NASIC Operational Analysis Center, where he served with distinction until August 2009, being awarded the Air Force Commendation Medal. In August 2009, Capt Buschelman again returned to the Air Force Institute of Technology to work toward the Ph.D. in Electrical Engineering, with an emphasis in laser radar, image processing, and nonparametric statistics.

Capt Buschelman has been selected for assignment at the Air Force Research Laboratory Munitions Directorate at Eglin AFB, FL, where he will continue to pursue research beneficial to the mission of the US Air Force. He has co-authored multiple technical publications in the fields of microwave engineering and laser radar technology.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

| | | | | | | |
|--|--------------------|---------------------|---------------------------------------|----------------------------|---|--|
| 1. REPORT DATE (DD-MM-YYYY) 27 Dec 2012 | | | 2. REPORT TYPE Dissertation | | 3. DATES COVERED (From — To) Aug 2009 — Dec 2012 | |
| 4. TITLE AND SUBTITLE A Nonparametric Approach to Segmentation of Ladar Images | | | | | 5a. CONTRACT NUMBER | |
| | | | | | 5b. GRANT NUMBER | |
| | | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) Eric A. Buschelman, Captain, USAF | | | | | 5d. PROJECT NUMBER N/A | |
| | | | | | 5e. TASK NUMBER | |
| | | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management 2950 Hobson Way WPAFB OH 45433-7765 | | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/DEE/ENG/12-07 | |
| 9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research Attn: Dr. Kent Miller (kent.miller@afosr.af.mil) 875 Randolph St Arlington, VA 22203 DSN: 426-8573 | | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR | |
| | | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION / AVAILABILITY STATEMENT Approval for public release; distribution is unlimited. | | | | | | |
| 13. SUPPLEMENTARY NOTES This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. | | | | | | |
| 14. ABSTRACT The advent of advanced laser radar (ladar) systems that record full-waveform signal data has inspired numerous inquiries which aspire to extract additional, previously unavailable, information about the illuminated scene from the collected data. The quality of the information, however, is often related to the limitations of the ladar camera used to collect the data. This research project uses full-waveform analysis of ladar signals, and basic principles of optics, to propose a new formulation for an accepted signal model. A new waveform model taking into account backscatter reflectance is the key to overcoming specific deficiencies of the ladar camera at hand, namely the ability to discern pulse-spreading effects of elongated targets. A concert of non-parametric statistics and familiar image processing methods are used to calculate the orientation angle of the illuminated objects, and the deficiency of the hardware is circumvented. Segmentation of the various ladar images performed as part of the angle estimation, and this is shown to be a new and effective strategy for analyzing the output of the AFIT ladar camera. | | | | | | |
| 15. SUBJECT TERMS ladar, laser radar, nonparametric, segmentation, classification, full-waveform analysis | | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON | |
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | Richard K. Martin | |
| U | U | U | UU | 160 | 19b. TELEPHONE NUMBER (include area code) 937-255-3636 x 4582, richard.martin@afit.af.mil | |